REVIEWS OF MODERN PHYSICS, VOLUME 94, APRIL–JUNE 2022

Security in quantum cryptography

Christopher Portmann*

Concordium AG, 6300 Zug, Switzerland

Renato Renner†

Institute for Theoretical Physics, ETH Zurich, 8093 Zurich, Switzerland

(published 29 June 2022)

Quantum cryptography exploits principles of quantum physics for the secure processing of information. A prominent example is secure communication, i.e., the task of transmitting confidential messages from one location to another. The cryptographic requirement here is that the transmitted messages remain inaccessible to anyone other than the designated recipients, even if the communication channel is untrustworthy. In classical cryptography, this can usually be guaranteed only under computational hardness assumptions, such as when factoring large integers is infeasible. In contrast, the security of quantum cryptography relies entirely on the laws of quantum mechanics. Here this physical notion of security is reviewed, with a focus on quantum key distribution and secure communication.

DOI: 10.1103/RevModPhys.94.025008

CONTENTS

*cp@concordium.com
†renner@ethz.ch

I. SECURITY FROM PHYSICAL PRINCIPLES

Communication theory is concerned with the task of making information available to different parties. The sender of a message x wants x to become accessible to a designated set of recipients. In *cryptography*, one adds to this a somewhat opposite requirement: that of restricting the availability of information. The sender of x also wants a guarantee that x will remain inaccessible to *adversaries*, i.e., parties other than the intended recipients. The term *security* refers to this additional guarantee.

Testing whether a communication protocol works correctly is easy. It suffices to compare the message x sent with the received one. Testing security, however, is more subtle. To ensure that an adversary cannot read x, one needs to exclude all physically possible eavesdropping strategies. Since there are infinitely many such strategies it is not possible, at least not by direct experiments, to prove that a cryptographic scheme is secure (although a successful hacking experiment would show the opposite).

But the situation is not as hopeless as this sounds. Security can be established indirectly, provided that one is ready to make certain assumptions about the capabilities of the adversaries. The weaker these assumptions are, the more confident we can be that they will apply to any realistic adversary, and hence that a cryptographic scheme based on them is actually secure.

The security of most cryptographic schemes used today relies on computational hardness assumptions. They correspond to constraints on the adversaries' computational resources. For example, it is assumed that adversaries do not have the capacity to factor large integers (Rivest, Shamir, and Adleman, 1978). This is a relatively strong assumption, justified merely by the belief that the currently known algorithms for factoring will not be substantially improved in the foreseeable future, and that quantum computers powerful enough to run Shor's efficient factoring algorithm (Shor, 1997) cannot be built. Cryptographic schemes whose security is based on assumptions of this type are commonly termed *computationally secure*.

In contrast, the main assumption that enters quantum cryptography is that adversaries are subject to the laws of quantum mechanics.[1] This assumption completely substitutes computational hardness assumptions; i.e., security holds even if the adversaries can use unbounded computational resources to process their information.[2] To distinguish this from computational security, the resulting security is sometimes termed *information theoretic*, reflecting the fact that it can be defined in terms of purely information-theoretic concepts (Shannon, 1949).

A. Completeness of quantum theory

The assumption that adversaries are subject to the laws of quantum mechanics appears to be rather straightforward to justify. Indeed, quantum mechanics is one of our best tested physical theories. To date no experiment has been able to detect deviations from its predictions. Of particular relevance to cryptography are nonclassical features of quantum mechanics, such as entanglement between remote subsystems, which have been tested by Bell experiments (Freedman and Clauser, 1972; Aspect, Grangier, and Roger, 1981; Aspect, Dalibard, and Roger, 1982; Tittel *et al.*, 1998; Weihs *et al.*, 1998; Rowe *et al.*, 2001; Christensen *et al.*, 2013; Giustina *et al.*, 2013, 2015; Hensen *et al.*, 2015; Shalm *et al.*, 2015; Rosenfeld *et al.*, 2017). However, the assumption that enters quantum cryptography concerns not only the *correctness* of quantum mechanics (as one may naively think) but also its *completeness*. This is an important point, and we therefore devote this entire section to it.

Quantum mechanics is a *nondeterministic* theory in the following sense. Even if we know, for instance, the polarization direction ψ of a photon to arbitrary accuracy, the theory will not in general allow us to predict with certainty the outcome z of a polarization measurement of, say, the vertical versus the horizontal direction. The statement that we can obtain from quantum mechanics may even be completely uninformative. For example, if the polarization ψ before the measurement were diagonal, the theory merely tells us that a measurement of the vertical versus the horizontal direction would yield both possible outcomes z with equal probability.

It is conceivable that nondeterminism is just a limitation of current quantum theory, rather than a fundamental property of nature. This would mean that there could exist another theory that would give better predictions. In the previous example, it could be that the photon, in addition to its polarization state ψ, has certain not yet discovered properties λ on which the measurement outcome z depends. A theory that takes into

[1]To prove security, one usually also requires that adversaries cannot manipulate the local devices (such as senders and receivers) of the legitimate parties. But this seemingly necessary requirement can be weakened: this is the topic of device-independent cryptography, which we discuss in Sec. IX.A.

[2]Although one may naturally also consider computationally secure quantum cryptography, which we do in Secs. VIII.F and IX.B.

account λ could then yield more informative predictions for z than quantum mechanics. If this were the case, then quantum mechanics could not be considered a complete theory.

Quantum cryptography is built on the use of physical systems, such as photons, as information carriers. The incompleteness of quantum mechanics would hence imply that the theory does not give a full account of all information contained in these systems. This would have severe consequences for security claims. For example, a cryptographic scheme for transmitting a confidential message x may be claimed to be secure on the grounds that the quantum state ψ of the information carriers gathered by an adversary is independent of x. Nonetheless, it could still be that the adversary's information carriers have an extra property λ that is not described by quantum theory and hence not included in ψ. The independence of ψ from x is then not sufficient to guarantee that the adversary could not learn the secret message.

A possible way around this problem is to simply assume that no adversary can access properties of physical systems, like λ in the previous example, that are not captured by their quantum state ψ. But such an assumption seems to be similarly difficult to justify as the nonexistence of an efficient factoring algorithm. The fact that we have not yet been able to discover λ does not mean that it does not exist (or that it cannot be discovered).

The problem can be resolved in a more fundamental manner. The solution is based on a long sequence of work dating back to that of Born (1926) and Einstein, Podolsky, and Rosen (1935), where the question regarding the completeness of quantum mechanics was raised. The central insight resulting from this work was that the set of possible theories that could improve the predictions of quantum mechanics is highly constrained. For example, no such theory can yield deterministic predictions, based on additional parameters λ, unless it is nonlocal[3] (Bell, 1964) and contextual (Bell, 1966; Kochen and Specker, 1967). More recently it was shown that no theory can improve the predictions of quantum mechanics unless it violates the requirement that measurement settings be chosen freely, i.e., independently of other parameters of the theory (Colbeck and Renner, 2011).[4] The completeness of quantum mechanics is hence implied by the assumption that physics does not prevent us from making free choices, an assumption that appears to be unavoidable in cryptography anyway (Ekert and Renner, 2014).

B. Correctness of quantum-theoretic description

In Sec. I.A we saw that the security of quantum cryptography crucially relies on the completeness of quantum mechanics, but that the latter can be derived from the requirement that one can make free choices. It is still necessary to assume that quantum mechanics is correct, in the sense that it accurately describes the hardware used for implementing a cryptographic protocol. But since quantum mechanics consists of a set of different rules, we should be more specific about what this correctness assumption really means.

Quantum cryptographic protocols are usually described within the framework of quantum information theory (Nielsen and Chuang, 2010), which provides the necessary formalism to talk about information carriers and operations on them. Any information carrier is modeled as a quantum system S with an associated Hilbert space $\mathcal{H}_S$, and the information encoded in S corresponds to its state. In the case of "classical" information, the different values x of a variable with range $\mathcal{X}$ are represented by different elements from a fixed orthonormal basis $\{|x\rangle\}_{x\in\mathcal{X}}$ of $\mathcal{H}_S$. If the marginal state of a system S has the form $\rho_S = \sum_x p_x|x\rangle\langle x|$, this means that S carries the value x with probability p_x. Any processing of information (including a measurement) corresponds to a change of the state of the involved information carriers and is represented mathematically by a trace-preserving, completely positive map.[5]

The modeling of real-world implementations in terms of these rather abstract information-theoretic notions is a highly nontrivial task. To illustrate this, take an optical scheme for quantum key distribution, where information is communicated by an encoding in the polarization of individual photons. This suggests a description where each photon sent over the optical channel is regarded as an individual quantum system. However, photons are just excitations of the electromagnetic field and thus *a priori* not objects with their own identity. (That is, they are indistinguishable.) A solution to this problem could be that one "labels" the photons by the time at which they are sent out; i.e., photons sent at different times are regarded as different quantum systems S. But there could be more than one photon emitted at a particular time, and these different photons either could or could not have the same frequency. One may now choose to take this into account by modeling the photon number and their frequency as internal degrees of freedoms of the system S. Or one could choose the frequency to be an additional system label so that photons with different frequencies would be regarded as different systems.

This example shows that the translation of an actual physical setup into the language of quantum information theory is prone to mistakes and is certainly not unique. Nonetheless, it is critical for security: if done incorrectly, the security statements, which are derived within quantum information theory, are vacuous. Particular care must be taken to ensure that no information carriers that are present in an implementation are omitted. A realistic photon source may sometimes emit two instead of only one photon whose polarization encodes the same value, and this second photon may be accessible to an eavesdropper (see Sec. IV.C). This possibility must therefore be included in the quantum information-theoretic description of that photon source. If it were not, it would represent a side channel to the adversary that is not accounted for by the security proof. Side channels may

[3]This concept is discussed in the context of device-independent cryptography in Sec. VI.D.

[4]More precisely, according to Bell and Aspect (2004) variables are "free" if they "have implications only in their future light cones." In other words, they are not correlated to anything outside their causal future. This notion has sometimes also been called "free will" (Conway and Kochen, 2006).

[5]We refer the interested reader to standard textbooks in quantum information theory, such as that of Nielsen and Chuang (2010), for a description of these concepts. An argument that justifies their use in the context of cryptography was given by Renes and Renner (2020).

also occur in other components, such as photon detectors; see the review of Scarani *et al.* (2009) for a general discussion of these practical aspects of quantum cryptography.

C. Overview of this review

In this review we focus on the information-theoretic layer of security proofs; i.e., we presume that we have a correct quantum information-theoretic description of the cryptographic hardware. The existence of such a description is indeed a standard assumption made for security proofs and is usually termed device dependence. It contrasts with device-independent cryptography, where this assumption is considerably relaxed; see Sec. VI.D for a discussion.

We start Sec. II by introducing general concepts from cryptography. From then on we focus largely on quantum key distribution (QKD), which currently is the most widespread application of quantum cryptography. It is also a concrete example with which to discuss security definitions, the underlying assumptions, and proof techniques. In Secs. VII and VIII we explain how these notions apply to cryptographic tasks other than key distribution.

II. CRYPTOGRAPHIC SECURITY DEFINITIONS

A. Real-world ideal-world paradigm

Cryptographic schemes are not usually perfectly secure. Rather, they provide a certain level of security that is quantified by one or several parameters. Take for instance an encryption scheme. It could be called perfectly secure only if we had a guarantee that an adversary could learn absolutely nothing about the encrypted message, something that turns out to be impossible to achieve in practice. Still, we have encryption schemes, such as those built in quantum cryptography, that are "almost perfectly secure." Thus, we need a quantitative definition that makes what this means precise.

Devising sensible quantitative definitions can be challenging. Consider a protocol that encrypts quantum information contained in a d-dimensional register A by applying a unitary U_k that depends on a uniformly chosen key $k \in \mathcal{K}$. It was proposed by Ambainis and Smith (2004), Hayden *et al.* (2004), and Dickinson and Nayak (2006) that the security of such a scheme may be defined by requiring that, for any state ρ_A,

$$\frac{1}{2}\left\| \frac{1}{|\mathcal{K}|}\sum_{k\in\mathcal{K}} U_k \rho_A U_k^\dagger - \tau_A \right\|_{\mathrm{tr}} \leq \varepsilon, \tag{1}$$

where $\varepsilon \geq 0$ is the security parameter, $\tau_A = (1/d)I$ is the fully mixed state, and $\|\cdot\|_{\mathrm{tr}}$ denotes the trace norm or Schatten 1-norm. The definition has been justified by the argument that an adversary who does not know the key k cannot distinguish the encryption of the state from τ_A (except with advantage[6] ε). However, it was later realized that this does not hide the information in the A system from an adversary who may hold a purification R of the information A (Ambainis, Bouda, and Winter, 2009). To take this into account, one should instead require that, for any ρ_{AR},

$$\frac{1}{2}\left\| \frac{1}{|\mathcal{K}|}\sum_{k\in\mathcal{K}} (U_k \otimes I_R)\rho_{AR}(U_k^\dagger \otimes I_R) - \tau_A \otimes \rho_R \right\|_{\mathrm{tr}} \leq \varepsilon,$$

where ρ_R is the reduced density operator of ρ_{AR}. Note that this criterion is not implied by Eq. (1) (Watrous, 2018).

As another example, early works on QKD (Mayers, 1996; Biham *et al.*, 2000; Shor and Preskill, 2000) measured the secrecy of a secret key in terms of the accessible information[7] between the key and all information that may be accessible to an adversary. In security proofs it was then shown that this value is small, apparently implying that the key is almost perfectly secret. Later one realizes, however, that the accessible information is not a good measure for secrecy: even if this measure is exponentially small in the key size, an adversary may be able to infer the second part of the key upon seeing the first part (König *et al.*, 2007). This makes the key unusable in many applications, such as encryption, as described in Sec. III.C.1.

Problems analogous to the previously outlined ones are well known in classical cryptography. They were addressed independently by Pfitzmann and Waidner (2000, 2001) and Canetti (2001), who built on a series of earlier works (Goldreich, Micali, and Wigderson, 1986; Beaver, 1992; Micali and Rogaway, 1992; Canetti, 2000) with a security paradigm that we refer to here as the "real-world ideal-world" paradigm. The gist of it lies in quantifying how well a real protocol for a cryptographic task can be distinguished from an ideal system that fulfills the task perfectly.

As a simple noncryptographic example, we consider channel coding, i.e., the task of constructing a noiseless channel from a noisy one. Suppose that Alice and Bob have access only to a noisy channel, as drawn in Fig. 1(b). To send a message, Alice will encode it in a larger message space that has redundancies. Upon reception, Bob will decode it, using the redundancies to correct errors (Nielsen and Chuang, 2010). Putting together the encoder, the noisy channel, and the decoder, as illustrated in Fig. 1(c), gives a new channel. Ideally, this constructed channel, which we call the real world, should behave like a perfect, noiseless channel [Fig. 1(a)], which we therefore call the ideal world. To quantify how well we achieved this goal, we measure how close the real world is to the ideal world.

For this, we consider a hypothetical game in which a distinguisher has black-box access to an unknown system, as shown in Fig. 2. The unknown system is, depending on a random bit B, either the real world ($B = 0$) or the ideal world ($B = 1$). The term black-box access means that the distinguisher is not provided with a description of the system, and, in particular, has no direct access to the bit B, but otherwise can interact arbitrarily with it. In the case of our noiseless channel construction problem, the distinguisher can generate

[6]See the text near Eq. (2) for a definition of the notion of a distinguishing advantage.

[7]This captures the information that a player may obtain by measuring her quantum state and is formally defined in Eq. (16) in Sec. III.C.1; see also Nielsen and Chuang (2010).

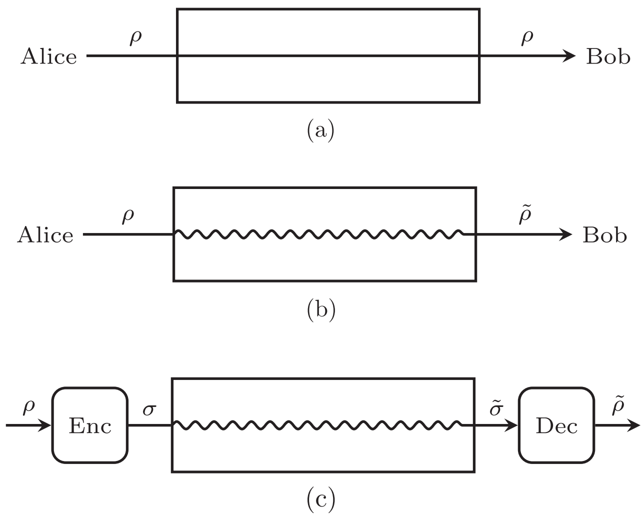

FIG. 1. (a) A noiseless channel that perfectly delivers the message from Alice to Bob. (b) A noisy channel that alters the message sent from Alice to Bob. (c) Alice encodes her message into a larger space, and Bob decodes it upon reception. In (a)—(c) each box represents a reactive system that produces an output upon receiving an input. Boxes with rounded corners are local operations performed by a party [such as encoding and decoding in (c)]. The rectangular box is a possibly noisy channel form Alice to Bob, which upon receiving Alice's input produces an output at Bob's end of the channel. The arrows represent quantum states transmitted from one system to another.

any joint state ρ_{AR} it desires, input the A part into the channel, and then measure the joint state of the channel output and its purification R. The distinguisher is then asked to guess whether it interacts with the real world ($B = 0$) or with the ideal one ($B = 1$). Let D be a random variable denoting the distinguisher's guess. The distinguishing advantage of the distinguisher is then defined as the difference between the probabilities that it guessed correctly or erroneously, namely,

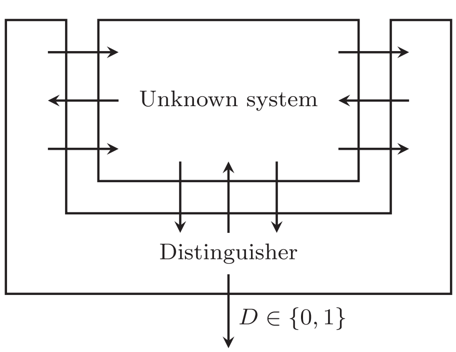

FIG. 2. In the real-world ideal-world paradigm, security is defined in terms of indistinguishability. A distinguisher has black-box access to a system that, depending on an unknown bit B, is either the real cryptographic protocol ($B = 0$) or an ideal functionality ($B = 1$). After interacting with the system, the distinguisher outputs a guess D for B. The real protocol is considered as secure as the ideal system if the success probability $\Pr[D = B]$ of the best possible distinguisher is close to that of a random guess, i.e., to $1/2$.

$$|\Pr[D = 0|B = 0] - \Pr[D = 0|B = 1]|. \tag{2}$$

The distinguishing advantage for a class of distinguishers (such as computationally bounded or unbounded distinguishers) is then defined as the supremum of Eq. (2) over all distinguishers in this set. For example, in the case of channel coding the distinguishing advantage for unbounded distinguishers corresponds to the diamond norm between the channels (Watrous, 2018). A protocol is considered secure if the distinguishing advantage is small, or, more accurately, the level of security of a protocol is parametrized by this advantage and the corresponding class of distinguishers.

In its essence, the real-world ideal-world paradigm avoids defining security; instead, it provides a simple description, the ideal world, of what should happen in the real world. In the example of channel coding, the real world might involve a complex noise model as well as encoding and decoding operations, whereas the ideal world is simply an identity map. When evaluating whether such a security statement is appropriate, one asks whether the ideal world captures what we need or if one should design a different ideal world.

A crucial property of the real-world ideal-world paradigm is that the resulting notion of security is composable. This means that the security of a protocol is guaranteed even if it is composed with other protocols to form a larger cryptographic system. In fact, to ensure composability, the notion of distinguishability has to be chosen appropriately. Specifically, the distinguisher must have access jointly to all information available normally to the honest parties as well as the adversary. The role of the distinguisher is hence to capture "the rest of the world," everything that exists around the system of interest. In particular, the distinguisher may choose the inputs to the protocol (which might come from a previously run protocol), receive its outputs (which could be used in a subsequent protocol), and simultaneously take the role of the adversary, possibly eavesdropping on the communication channels and tampering with messages.

B. The abstract cryptography framework

In modern cryptography, security claims and their proofs are usually phrased within a theoretical framework. The framework not only provides a common language but also ensures composability, in the previously described sense. That is, security claims that hold for individual components can be turned into a security claim for the more complex cryptographic scheme built from them. The first frameworks to achieve this for classical cryptography were the reactive simulatability framework of Pfitzmann and Waidner (2000, 2001) and the universal composability framework of Canetti (2001), which both used the real-world ideal-world paradigm.

These frameworks have been further developed (Backes, Pfitzmann, and Waidner, 2004, 2007; Canetti *et al.*, 2007; Canetti, 2020) and several variations have been proposed (Mateus, Mitchell, and Scedrov, 2003; Canetti *et al.*, 2006a, 2006b; Küsters, 2006; Mitchell *et al.*, 2006; Hofheinz and Shoup, 2015). The differences between them concern mostly how they describe information-processing systems, i.e., how the individual devices carry out computations and how they schedule messages when communicating (e.g., synchronously or asynchronously). While this modeling was based mostly on classical notions of computation and communication, the frameworks were also adapted to quantum cryptography by Ben-Or and Mayers (2004) and Unruh (2004, 2010).

Maurer and Renner (2011) [see also Maurer (2012) and Maurer and Renner (2016)] proposed a framework, abstract cryptography (AC), that is largely independent of the underlying modeling of the information-processing devices, and therefore applies equally to classical and quantum settings. We use it for the presentation here, for it enables a self-contained description without the need to specify unnecessary technical details.[8] In the following, we describe the two basic paradigms on which the framework is based, abstraction and constructibility.

1. Abstraction

The traditional approach to defining security (used in all the frameworks cited except for AC) can be seen as bottom up. One first defines (at a low level) a computational model (such as a Turing machine) and then proceeds by modeling how the machines communicate (such as by writing to and reading from shared tapes). Next one introduces higher-level notions such as indistinguishability. Finally, these notions are used to define security.

In contrast, AC uses a top-down approach. To state definitions and develop a theory, one starts at the other end, the highest possible level of abstraction. There cryptographic systems are simply regarded as elements of a set that can be combined to form new systems. One then proceeds down to lower levels of abstraction, introducing in each of them the minimum necessary specializations. Only on these lower levels is it modeled how exactly the cryptographic systems process information and how they communicate when they are combined (e.g., synchronously or asynchronously). The notion of indistinguishability is first defined on the highest abstraction level as an arbitrary metric on the set of cryptographic systems. On lower abstraction levels it can then be instantiated in different ways, such as to capture the distinguishing power of a computationally bounded or unbounded environment.

Abstraction not only has the advantage that it generalizes the treatment but usually also simplifies it, as unnecessary specificities are avoided. It may be compared to the use of group theory in mathematics, which is an abstraction of more special concepts such as matrix multiplication. In a bottom-up approach, one would start by introducing a rule for taking the product between matrices and then, based on that rule, studying the properties of the multiplication operation. In contrast to this, the top-down approach taken here corresponds to first defining the abstract multiplication group and then proving theorems that are already on this level.

[8]The security of QKD could equivalently be modeled based on the work of Unruh (2010), with minor adaptations to capture finite statements instead of only asymptotics.

2. Constructibility

Cryptography can be regarded as a resource theory, where certain desired resources are constructed from a set of given resources.[9] The constructions are defined by protocols. For example, a QKD protocol uses a quantum communication channel together with an authentic channel[10] as resources to construct the resource of a secret key. The latter resource may then be used by other protocols, such as an encryption protocol, to construct a secure communication channel, which is again a resource. Similarly, the authentic channel used by the QKD protocol can itself be constructed from an insecure channel resource and short uniform secret key (Wegman and Carter, 1981). And given a weak secret key (i.e., not necessarily uniform and not perfectly correlated randomness shared by the communication partners) and two-way insecure channels, one may construct an almost perfect secret key (i.e., uniform and perfectly correlated randomness) using so-called nonmalleable extractors (Renner and Wolf, 2003; Dodis and Wichs, 2009; Aggarwal *et al.*, 2019). Composing the authentication protocol with the QKD protocol results in a scheme that constructs a long secret key from a short secret key and insecure channels, and composing this again with nonmalleable extractors constructs this long key from only a weak key and insecure channels. Part of the resulting long secret key can be used in further rounds of authentication and QKD to produce an even more secret key. This is illustrated in Fig. 3 and discussed in Sec. VII.

The resources used and constructed in cryptography are interactive systems shared between players. A system that distributes a secret key or the different types of channels mentioned previously are examples of such resources. These are formalized on an abstract level in Sec. II.D, and possible instantiations are discussed in Sec. II.G. Static resources such as coherent states (Baumgratz, Cramer, and Plenio, 2014; Ma *et al.*, 2019) can be seen as a special case of these.

C. Example: The one-time pad

In this section, we describe how the previously introduced notions are employed to specify the security of a cryptographic protocol. For this we consider a concrete example, one-time pad (OTP) encryption (Vernam, 1926). The OTP assumes that the players have access to an authentic channel, i.e., one that provides the receiver with the guarantee that the messages received come from the correct sender, but there is no guarantee about the secrecy of the messages sent on such a channel; i.e., they may leak to Eve. The OTP also requires the players to have access to a secret key. These two resources are drawn as boxes with square corners in Fig. 4. According to the

[9]This view is widespread in cryptography and made formal by composable frameworks (Pfitzmann and Waidner, 2000, 2001; Canetti, 2001; Maurer and Renner, 2011). Resource theories have also been used in many different ways to capture certain aspects of quantum mechanics; see the reviews by Streltsov, Adesso, and Plenio (2017) and Chitambar and Gour (2019).

[10]An authentic channel guarantees that the message received comes from a legitimate sender and has not been tampered with or generated by an adversary.

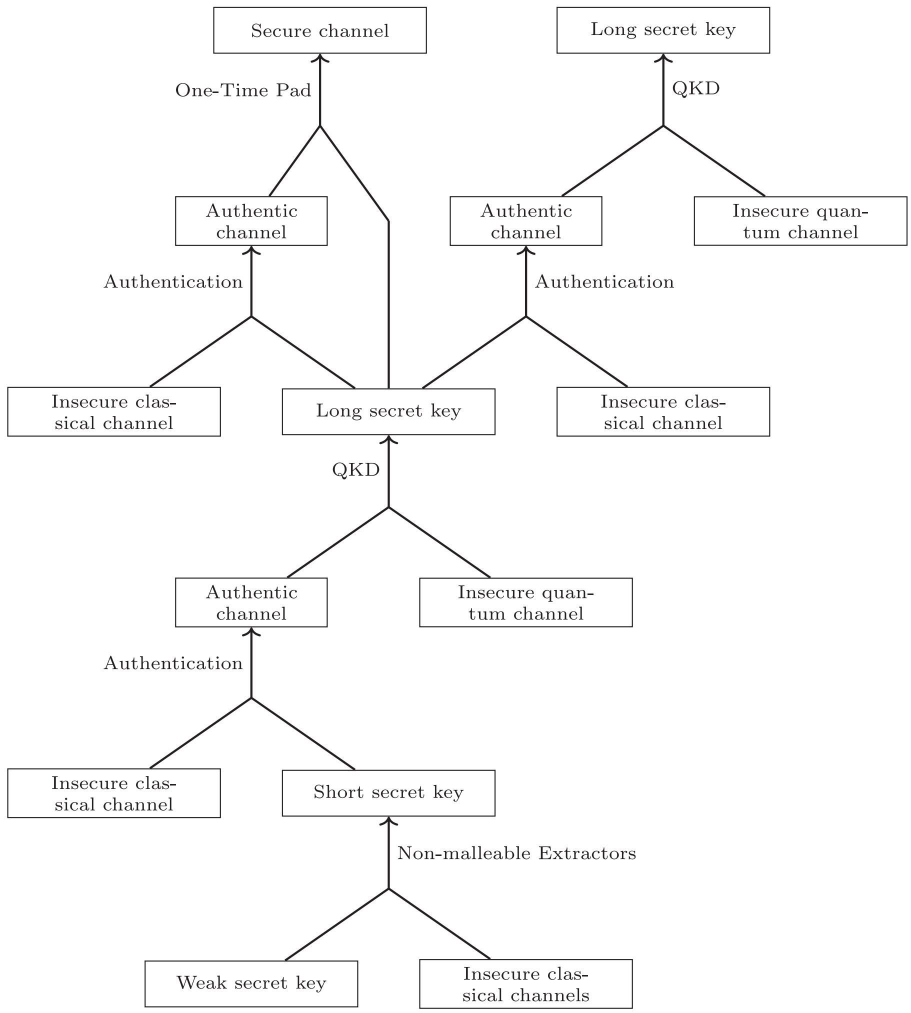

FIG. 3. A constructive view of cryptography. A cryptographic protocol uses weak resources to construct other (stronger) resources. These resources are depicted in the boxes, and the arrows are protocols. Each box is a one-time-use resource, so the same resource appears in multiple boxes if different protocols require it. The long secret key resource in the center is split in three shorter keys, each of which is used by a separate protocol. The example of secure message transmission illustrated here is discussed in Sec. VII.

protocol, the sender Alice encrypts a message x as $y \coloneqq x \oplus k$, where k is the secret key and $\oplus$ denotes the bitwise exclusive OR operation. The ciphertext y is then sent over an authentic channel to the receiver Bob, who decrypts it by carrying out the operation $x = y \oplus k$. At the same time, y may also leak to an adversary Eve.

In this example, the goal of the OTP is to add confidentiality to an authentic channel[11]; i.e., the ideal system is a secure channel, drawn as a box with square corners in Fig. 4(b). This is a channel that leaks the message size but no other information to Eve. It is straightforward to verify that in the real system, provided that the key k is uniformly distributed over bit strings of the same length as the message x, the ciphertext y is statistically independent of the message x. The ciphertext y hence does not provide Eve with any information about x, except potentially for its length $|x|$. It thus constructs a secure channel from Alice to Bob.

To make the real and ideal systems comparable, we consider an entire class of systems, which are obtained by appending an arbitrary system, called a simulator, to the Eve interface of the ideal channel resource. The systems from this class are sometimes called relaxations (of the ideal system). The idea is that none of these relaxations can be more useful to Eve than the original ideal channel, because she may herself always carry out the task of the simulator. Security now means that the real system is indistinguishable from at least one relaxation of the ideal system. In our example of the OTP, such a relaxation may be obtained by a simulator that simply generates a random string of length $|x|$ and outputs it at the Eve interface, as depicted in Fig. 4(b).

To establish security of OTP encryption, it is therefore sufficient to show that the real system depicted in Fig. 4(a) is indistinguishable from the relaxation of the ideal secure channel shown in Fig. 4(b). That is, the two systems must behave identically when they interact with a distinguisher.

[11]Alternatively, one may use the OTP with a completely insecure channel, and thus obtain a malleable confidential channel (Maurer, Rüedlinger, and Tackmann, 2012).

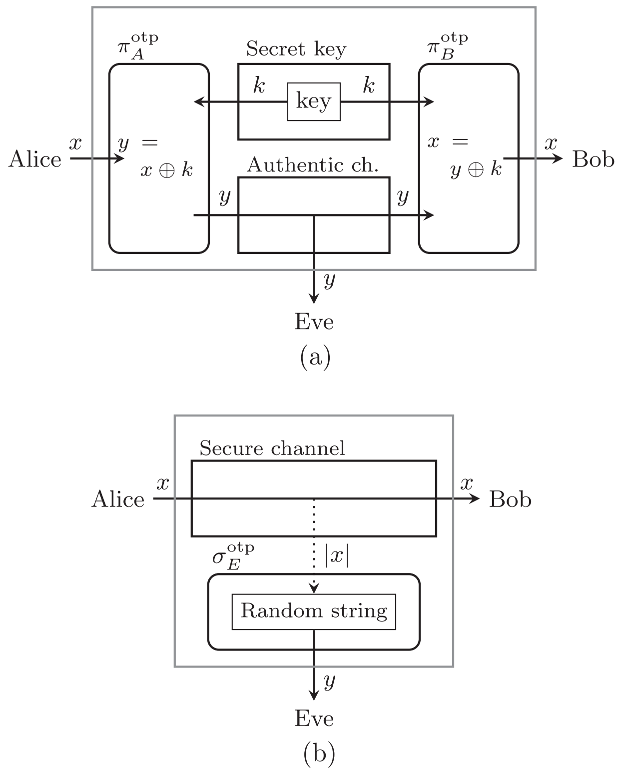

FIG. 4. Real and ideal one-time pad systems. Boxes with rounded corners are local systems executed at Alice's, Bob's, or Eve's interfaces. The rectangular boxes are shared resources modeling channels or shared keys. Arrows represent the transmission of messages between systems or to the environment (distinguisher). The real world is depicted in (a). The protocol consists of a part $\pi_A^{\rm OTP}$ executed by Alice (who has access to the interfaces on the left-hand side) and a part $\pi_B^{\rm OTP}$ executed by Bob (on the right-hand side). It takes a message x at Alice's outer interface as well as a key k and outputs a ciphertext y toward the authentic channel. Bob's part of the protocol takes y and k as input and outputs the decrypted message. The channel may leak y at Eve's interface (at the bottom). The ideal world is depicted in (b). The secure channel transmits the message perfectly from Alice's to Bob's interface, leaking only the message length at Eve's interface. The simulator $\sigma_E^{\rm OTP}$ generates a random string y of length $|x|$, making the real and ideal systems perfectly indistinguishable. (a) The real OTP system consists of the OTP protocol $(\pi_A^{\rm OTP}, \pi_B^{\rm OTP})$ together with a secret key and authentic channel resources. (b) The ideal OTP system consists of the ideal secure channel and a simulator $\sigma_E^{\rm OTP}$.

This is indeed the case. For both of them, if the distinguisher inputs x at Alice's interface, the same string x is output at Bob's interface and a uniformly random string of length $|x|$ is output at Eve's interface. The two systems are thus perfectly indistinguishable: if the distinguisher were to take a guess at which of the two it is interacting with, it would be correct with probability exactly $1/2$. In this sense, the OTP construction is perfectly secure.

If two systems are indistinguishable, they can be used interchangeably in any setting. For example, let some protocol π' be proven secure if Alice and Bob are connected by a secure channel. Since the OTP constructs such a channel, it can be used in lieu of the secure channel and composed with π'. Or, equivalently, the contrapositive: if composing the OTP and π' were to leak some vital information, which would not happen with a secure channel, a distinguisher that is given either the real or the ideal system could run π' internally and check to see whether this leak occurs to find out with which of the two it is interacting.

D. Abstract theory of cryptographic systems

We previously introduced the concepts of resources, protocols, and simulators in an informal manner. Now, following the spirit of the AC framework described in Sec. II.B, we provide an axiomatic specification of these concepts. This will allow us to give a definition of cryptographic security, which is precise, but at the same time largely independent of implementation details. In particular, it does not depend on the underlying computational model or the scheduling of messages exchanged between the systems.

While this abstract approach to defining security is rather universal, we note that, when describing concrete systems and their compositions such as those depicted in Fig. 4, their behavior must be specified in detail. This may be done using various frameworks for modeling interactive quantum systems such as the quantum combs of Chiribella, D'Ariano, and Perinotti (2009) and the causal boxes of Portmann *et al.* (2017). This is discussed further in Sec. II.G.

Nevertheless, the definitions that follow refer to an abstract notion of a system. Following the previously mentioned idea of abstraction and continuing the analogy to group theory used in Sec. II.B, it is sufficient to think of systems as objects on which certain operations are defined, such as their composition. We now consider two types of systems, which we call resources and converters and which have slightly different properties.

1. Resources

A resource is a system with interfaces specified by a set $\mathcal{I}$ (such as $\mathcal{I} = \{A, B, E\}$). Each interface $i \in \mathcal{I}$ models how a player i can access the system (e.g., how it can provide inputs and read outputs). Examples of resources are a communication channel or any of the objects that appear in Fig. 3 as a box. We sometimes use the term $\mathcal{I}$ resource to specify the interface set. Resources are equipped with a parallel composition operator, denoted by $\|$, that maps two $\mathcal{I}$ resources to another $\mathcal{I}$ resource.

2. Converters

A converter is a system with two interfaces, an inside interface and an outside interface. A converter can be appended to a resource, converting it into a new resource. For this the inside interface connects to an interface of a resource and the outside interface becomes the new interface of the new resource; see the OTP example in Fig. 4, where the gray boxes are new resources resulting from composing resources and converters. We write either $\alpha_i \mathcal{R}$ or $\mathcal{R}\alpha_i$ to denote the new resource with the converter α connected at the

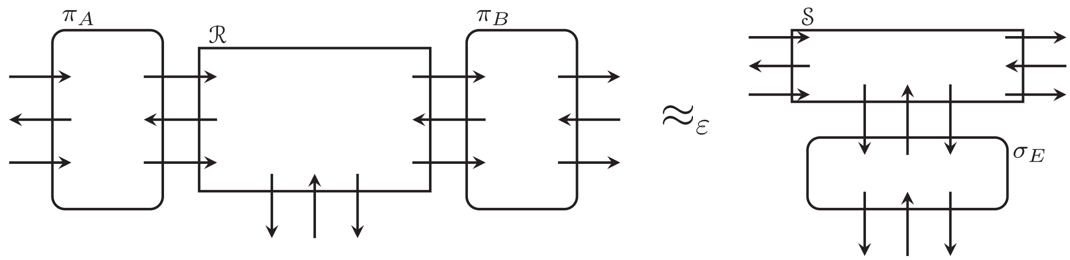

FIG. 5. Depiction of Definition 1. A protocol (π_A, π_B) constructs $\mathcal{S}$ from $\mathcal{R}$ within ε if the condition illustrated here holds. The sequences of arrows at the interfaces between the objects represent arbitrary rounds of communication.

interface i of $\mathcal{R}$.[12] Simulators and protocols are examples of converters that are discussed later.

Converters can be composed of themselves. There are two ways of doing this, referred to as serial and parallel composition. These are defined as

$$(\alpha\beta)_i\mathcal{R} := \alpha_i(\beta_i\mathcal{R})$$

and

$$(\alpha\|\beta)_i(\mathcal{R}\|\mathcal{S}) := (\alpha_i\mathcal{R})\|(\beta_i\mathcal{S}),$$

respectively.

3. Protocols

A cryptographic protocol is a family $\alpha = \{\alpha_i\}_i$ of converters (one for every honest player). A protocol can be applied to a resource $\mathcal{R}$, giving a new resource denoted by $\alpha\mathcal{R}$ or $\mathcal{R}\alpha$. This resource is obtained by connecting each member of the family to the interface specified by its index.

4. Metric

As explained in Sec. II.A, the distance between resources can be quantified using the notion of distinguishers. More generally, one may in principle consider any arbitrary pseudometric $d(\cdot,\cdot)$ such that the following conditions hold[13]:

$$\begin{aligned} &\text{(identity)} && d(\mathcal{R},\mathcal{R}) = 0,\\ &\text{(symmetry)} && d(\mathcal{R},\mathcal{S}) = d(\mathcal{S},\mathcal{R}),\\ &\text{(triangle inequality)} && d(\mathcal{R},\mathcal{S}) \le d(\mathcal{R},\mathcal{T}) + d(\mathcal{T},\mathcal{S}).\end{aligned}$$

Furthermore, the pseudometric must be nonincreasing under composition with resources and converters.[14] This means that, for any converter α and resources $\mathcal{R}$, $\mathcal{S}$, and $\mathcal{T}$,

$$d(\alpha\mathcal{R},\alpha\mathcal{S}) \le d(\mathcal{R},\mathcal{S}), \qquad d(\mathcal{R}\|\mathcal{T},\mathcal{S}\|\mathcal{T}) \le d(\mathcal{R},\mathcal{S}).$$

[12]There is no mathematical difference between $\alpha_i\mathcal{R}$ and $\mathcal{R}\alpha_i$. It sometimes simplifies the notation to have the converters for some players written on the right of the resource and the ones for others on the left, rather than having all of them at the same side; hence the two notations.

[13]If $d(\mathcal{R},\mathcal{S}) = 0 \Rightarrow \mathcal{R} = \mathcal{S}$ also holds, then d is a metric.

[14]This holds only for information-theoretic security, which is the topic of most of this review.

In this review we often simply write $\mathcal{R} \approx_\varepsilon \mathcal{S}$ instead of $d(\mathcal{R},\mathcal{S}) \le \varepsilon$.

E. Security definition

We are now ready to define the security of a cryptographic protocol. We do so in Definition 1 in the three-player setting, for honest Alice and Bob and dishonest Eve, and illustrate this definition in Fig. 5. Thus, in the following all resources have three interfaces, denoted A, B, and E, and we consider honest behaviors [given by a protocol (π_A, π_B)] at the A and B interfaces, but arbitrary behavior at the E interface. See Maurer and Renner (2011) for the general case, in which arbitrary players can be dishonest.

Definition 1 [cryptographic security (Maurer and Renner, 2011)].—Let $\pi_{AB} = (\pi_A, \pi_B)$ be a protocol and $\mathcal{R}$ and $\mathcal{S}$ be two resources. We say that π_{AB} constructs $\mathcal{S}$ from $\mathcal{R}$ within ε, denoted by

$$\mathcal{R} \xrightarrow{\pi,\varepsilon} \mathcal{S},$$

if there is a converter σ_E (called a *simulator*) such that

$$d(\pi_{AB}\mathcal{R}, \mathcal{S}\sigma_E) \le \varepsilon. \tag{3}$$

If it is clear from the context which resources $\mathcal{R}$ and $\mathcal{S}$ are meant, we simply say that π_{AB} is ε secure.

Although this security definition does not refer to any computational notions, one usually considers only those protocols whose converters are computationally efficient.[15] Furthermore, if one requires security to hold under composition with protocols that have only computational security, it is necessary to restrict the choice of the simulator σ_E to converters that are computationally efficient. All the converters and resources considered in this work are efficient in the standard sense, so we make no further mention of this.

For a given protocol, we usually want to make several security statements, such as one about what is achieved in the presence of an adversary (sometimes referred to as either the soundness or the security of a protocol) and another about what is achieved when no adversary is present (usually called either

[15]In principle, any reasonable notion of efficiency could be considered here. However, if one takes the common asymptotic notion of computational complexity classes, one would need to describe systems in terms of a computational model that enables such asymptotic considerations.

completeness or correctness[16]). These two cases are captured by considering different resources $\mathcal{R}$ and $\mathcal{S}$, but the same protocol π_{AB}. We illustrate this in Sec. III for the case of QKD.

If two protocols π and π' are ε and ε' secure, the composition of the two is $\varepsilon+\varepsilon'$ secure. More precisely, let protocols π and π' construct $\mathcal{S}$ from $\mathcal{R}$ and $\mathcal{T}$ from $\mathcal{S}$ within ε and ε', respectively, i.e.,

$$\mathcal{R} \xrightarrow{\pi,\varepsilon} \mathcal{S}, \qquad \mathcal{S} \xrightarrow{\pi',\varepsilon'} \mathcal{T}.$$

It is then a consequence of the triangle inequality of the distinguishing metric that $\pi'\pi$ constructs $\mathcal{T}$ from $\mathcal{R}$ within $\varepsilon+\varepsilon'$ as follows:

$$\mathcal{R} \xrightarrow{\pi'\pi,\varepsilon+\varepsilon'} \mathcal{T}.$$

A similar statement holds for parallel composition. Let π and π' construct $\mathcal{S}$ and $\mathcal{S}'$ from $\mathcal{R}$ and $\mathcal{R}'$ within ε and ε', respectively, i.e.,

$$\mathcal{R} \xrightarrow{\pi,\varepsilon} \mathcal{S}, \qquad \mathcal{R}' \xrightarrow{\pi',\varepsilon'} \mathcal{S}'.$$

If these resources and protocols are composed in parallel, we find that $\pi\|\pi'$ constructs $\mathcal{S}\|\mathcal{S}'$ from $\mathcal{R}\|\mathcal{R}'$ within $\varepsilon+\varepsilon'$ as follows:

$$\mathcal{R}\|\mathcal{R}' \to \pi\|\pi', \varepsilon + \varepsilon'\mathcal{S}\|\mathcal{S}'.$$

Proofs of these statements were given by Maurer and Renner (2011) and Maurer (2012).

F. Interpretation of the security parameter

Any pseudometric that satisfies the basic axioms can be used in Definition 1. However, the usual pseudometric, which was introduced in Eq. (2) in Sec. II.A, is the distinguishing advantage. For two resources $\mathcal{R}$ and $\mathcal{S}$ and a distinguisher $\mathfrak{D}$, Eq. (2) may be rewritten as

$$d^{\mathfrak{D}}(\mathcal{R},\mathcal{S}) := |\Pr[\mathfrak{D}(\mathcal{R})=0] - \Pr[\mathfrak{D}(\mathcal{S})=0]|, \qquad (4)$$

where $\mathfrak{D}(\mathcal{R})$ and $\mathfrak{D}(\mathcal{S})$ are the random variables corresponding to the output of the distinguisher when they interact with $\mathcal{R}$ and $\mathcal{S}$, respectively. Alternatively, one may define the distinguishing advantage for $\mathfrak{D}$ as

$$d^{\mathfrak{D}}(\mathcal{R},\mathcal{S}) := |2p^{\mathfrak{D}}_{\text{distinguish}}(\mathcal{R},\mathcal{S}) - 1|, \qquad (5)$$

where $p^{\mathfrak{D}}_{\text{distinguish}}(\mathcal{R},\mathcal{S})$ is the probability of $\mathfrak{D}$ correctly guessing with which of $\mathcal{R}$ or $\mathcal{S}$ it is interacting when either one is chosen with probability $1/2$, i.e.,

$$p^{\mathfrak{D}}_{\text{distinguish}}(\mathcal{R},\mathcal{S}) := \tfrac{1}{2}\Pr[\mathfrak{D}(\mathcal{R})=0] + \tfrac{1}{2}\Pr[\mathfrak{D}(\mathcal{S})=1].$$

Equations (4) and (5) are equivalent.

One then takes the supremum of this expression over all distinguishers $\mathfrak{D}$ of a given class $\mathbb{D}$, i.e.,

$$d^{\mathbb{D}}(\mathcal{R},\mathcal{S}) := \sup_{\mathfrak{D}\in\mathbb{D}} d^{\mathfrak{D}}(\mathcal{R},\mathcal{S}). \qquad (6)$$

The class $\mathbb{D}$ may be restricted to a particular set of systems (such as those that are computationally efficient). The strongest security notion corresponds to not imposing any restriction on the set of distinguishers (beyond what is allowed by physical laws), which is the one considered in most of this review and which we denote as

$$d(\mathcal{R},\mathcal{S}) \le \varepsilon \quad \text{or} \quad \mathcal{R} \approx_\varepsilon \mathcal{S}.$$

The distinguishing advantage is of particular importance because it has an operational interpretation. If the distinguisher notices a difference between the two, then something in the real setting did not behave ideally. This can be loosely interpreted as a failure occurring. If a distinguisher can guess correctly with probability 1 with which system it is interacting [i.e., $p_{\text{distinguish}}(\mathcal{R},\mathcal{S})=1$], a failure must occur systematically. If, conversely, it can only guess correctly with probability $1/2$ (which corresponds to a random guess), this means that the real system always behaves like the ideal one, and hence no failure occurs at all. The practically relevant cases are those in between. As shown in Appendix A, a guessing probability $p_{\text{distinguish}}(\mathcal{R},\mathcal{S})=p$ corresponds to a failure with probability $\varepsilon = 2p-1$, which is exactly the distinguishing advantage. The latter can thus be interpreted as the probability that a failure occurs in the real protocol. This operational interpretation is crucial for applications, where one must be able to specify what maximum value ε one is ready to tolerate.

A bound on the security ε of a protocol does not, however, tell us how "bad" this failure is. For example, a key distribution protocol that produces perfectly uniform keys for Alice and Bob (but with probability ε the keys of Alice and Bob are different) is ε secure. Likewise, a protocol that gives 1 bit of the key to Eve with probability ε (but is perfect otherwise) and another protocol that gives the entire key to Eve with probability ε (but is perfect otherwise) are both ε secure as well. One could argue that leaking the entire key is worse than leaking 1 bit, which is worse than not leaking anything but generating mismatching keys, and this should be reflected in the level of security of the protocol. However, leaking 1 bit can be as bad as leaking the entire key if only 1 bit of the message is vital, and this happens to be the bit obtained by Eve. Having mismatching keys and therefore misinterpreting a message could have more dire consequences than leaking the message to Eve. How bad a failure this is depends on the use of the protocol and, since the purpose of cryptographic security is to make a security statement that is valid for all contexts, bounding the probability that a failure (grave or not) occurs is the best it can do.

This is particularly relevant if one considers larger cryptographic tasks that may, for instance, use the key distribution numerous times as a subprotocol. Since, as described, a

[16]In the QKD literature, correctness has another meaning: it captures the property that Alice and Bob end up with identical keys when Eve is active. The term robustness is traditionally used in the QKD literature to denote the performance of a QKD protocol under honest (noisy) conditions; see Sec. III.B.4 for a discussion of the relation between completeness and robustness.

security bound gives one no idea of the gravity of a failure, the failure of the key distribution protocol could have an impact on the entire cryptographic system. For example, if the key is used to authenticate later communication, the security of the latter may be affected by a failure in key distribution. This makes it necessary to choose the probability ε of a failure in any protocol to be small enough that the accumulation of all possible failure probabilities used for the larger cryptographic task are still small. One way of doing this is to increase the security parameter of a protocol on a regular basis: e.g., once a year the parameters are tweaked such that the new probability of a failure is divided by 2. If the accumulated failure during the first year is given by ε, then the total failure over an arbitrarily long lifetime of the system is bounded by $2\varepsilon > \varepsilon + \varepsilon/2 + \varepsilon/4 + \cdots$.

G. Instantiating systems

As previously mentioned, specifying a concrete behavior of a system requires a model of systems that satisfies the axioms presented in Sec. II.D, i.e., provides composition and a pseudometric with the required properties. In most of this review we consider interactive quantum systems with sequential scheduling; i.e., a system receives a quantum message, then sends a quantum message, then receives a quantum message, etc. Such systems were analyzed independently by Gutoski and Watrous (2007), Chiribella, D'Ariano, and Perinotti (2009), Hardy (2011, 2012, 2015), and Gutoski (2012) [see also Hardy (2005) and Hardy (2007)], to which we refer in the following using the term from Chiribella, D'Ariano, and Perinotti (2009), namely, quantum combs. Quantum combs are a generalization of random systems (Maurer, 2002; Maurer, Pietrzak, and Renner, 2007) to quantum information theory.

What these works essentially show is that an interactive system that receives the ith input in register A_i and produces the ith output in register B_i and which processes n inputs can be fully described by a completely positive, trace-preserving (CPTP) map

$$\mathcal{E}:\mathcal{L}\left(\bigotimes_{i=1}^{n}\mathcal{H}_{A_i}\right)\to\mathcal{L}\left(\bigotimes_{i=1}^{n}\mathcal{H}_{B_i}\right).$$

Conversely, any such CPTP map corresponds to an interactive system if it respects causality; i.e., if for any $j \leq n$ and any $\rho,\sigma \in \mathcal{L}(\bigotimes_{i=1}^{n}\mathcal{H}_{A_i})$ with $\mathrm{tr}_{A_{>j}}(\rho) = \mathrm{tr}_{A_{>j}}\sigma$ we have

$$\mathrm{tr}_{B_{>j}}[\mathcal{E}(\rho)] = \mathrm{tr}_{B_{>j}}[\mathcal{E}(\sigma)],$$

where $X_{>j} \coloneqq \bigotimes_{i=j+1}^{n} X_i$.

Systems such as the resources and converters for the one-time pad in Fig. 4 (or the quantum key distribution systems that come in Sec. III) all correspond to specific quantum combs. (Nonetheless, we usually give informal descriptions of such systems rather than using the comb formalism, especially when the details of their behavior are not relevant to our claims.) The only results discussed in this review that cannot be modeled as quantum combs are the relativistic systems reviewed in Sec. VIII.E, which require a more complex model of systems that can capture space-time and also satisfy the required axioms, such as the causal boxes of Portmann *et al.* (2017).

III. DEFINING SECURITY OF QKD

The first QKD protocols were proposed independently by Bennett and Brassard (1984), inspired by early work on quantum money by Wiesner (1983), and by Ekert (1991). The original papers discussed security in the presence of an eavesdropper that could perform only limited operations on the quantum channel. The models of security evolved over time (a review of these is given in Sec. III.C) and the security criterion used today was introduced in 2005 (Ben-Or *et al.*, 2005; Renner, 2005; Renner and König, 2005), the so-called trace distance criterion. It was argued as follows that ρ_{KE}, the joint state of the final key K and the quantum information gathered by an eavesdropper E, must be close to an ideal key τ_K that is perfectly uniform and independent from the adversary's information ρ_E:

$$(1 - p_{\mathrm{abort}})D(\rho_{KE}, \tau_K \otimes \rho_E) \leq \varepsilon, \tag{7}$$

where p_{abort} is the probability that the protocol aborts,[17] $D(\cdot,\cdot)$ is the trace distance,[18] and $\varepsilon \in [0, 1]$ is a small real number. This security criterion was discussed within the cryptography frameworks introduced in Sec. II by Ben-Or *et al.* (2005) and Müller-Quade and Renner (2009); see also Appendix A.

We note that Eq. (7) captures only how much an adversary knows about the key (called secrecy in the QKD literature). A QKD scheme must additionally guarantee that Alice and Bob hold the same key with high probability (called correctness), and that under reasonably noisy conditions a QKD scheme produces a key with high probability (called robustness). In this section, we describe how these security notions fit into the general framework described in Sec. II. For this we first explain in Sec. III.A how to use the AC framework to model the task achieved by a QKD protocol (namely, constructing a secret key resource from an insecure quantum channel and an authentic classical channel) and write out the corresponding security definitions. In Sec. III.B we then show how to derive secrecy, correctness, and robustness from these security definitions. And finally in Sec. III.C we review other security definitions that have appeared in the literature and explain how they relate to the trace distance criterion, namely, Eq. (7).

A. The real and ideal QKD systems

To apply the general AC security definition to QKD, we first need to specify the ideal key resource, which we do in Sec. III.A.1. Likewise, we specify in Sec. III.A.2 the real QKD system consisting of the protocol, an authentic classical

[17] Renner (2005) introduced Eq. (7) with a subnormalized state ρ_{KE}, with $\mathrm{tr}(\rho_{KE}) = 1 - p_{\mathrm{abort}}$, instead of explicitly writing the factor $1 - p_{\mathrm{abort}}$. The two formulations, however, are mathematically equivalent.

[18] This metric corresponds to the distinguishing advantage between two quantum states and is formally defined in Appendix A.

channel, and an insecure quantum channel. Plugging these systems into Definition 1, we obtain in Sec. III.A.3 the security criteria for QKD.

1. Ideal key

The goal of a key distribution protocol is to generate a secret key shared between two players Alice and Bob. One can represent such a resource by a box, one end of which is in Alice's lab and another in Bob's. It provides each of them with a secret key of a given length but does not give Eve any information about the key. This is illustrated in Fig. 6(a), and is the key resource that we used in the OTP construction [Fig. 4(a)].

However, if we want to realize such a functionality with QKD, there is a caveat: an eavesdropper can always prevent any real QKD protocol from generating a key by cutting or jumbling the communication lines between Alice and Bob, and this must be taken into account in the definition of the ideal resource. This box thus also has an interface accessible to Eve, which provides her with a switch that, when pressed, prevents the box from generating this key. We depict this in Fig. 6(b).

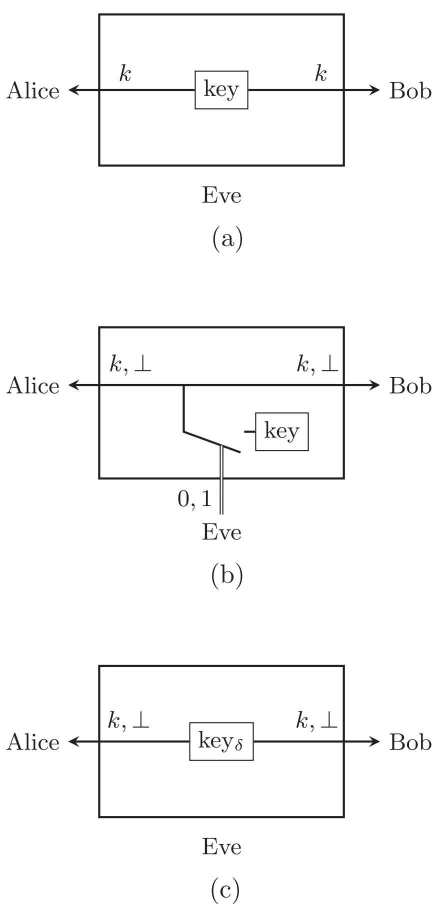

FIG. 6. Depictions of some shared secret key resources. (a) A resource that always gives a key k to Alice and Bob, and nothing to Eve. (b) A resource that allows Eve to decide whether Alice and Bob get a key k or an error $\perp$. (c) A resource that generates a perfect key with probability $1-\delta$ and outputs an error $\perp$ with probability δ.

If a OTP protocol uses the key generated by the resource of Fig. 6(b), we need to consider two cases. If Eve prevents a key from being generated, the construction is trivially secure: in this case, Alice and Bob do not have a key and therefore cannot send any messages. And in the case where a key is generated we have the situation depicted by Fig. 6(a), which is the situation that we already analyzed in Sec. II.C.

As previously explained, an adversary can prevent a key from getting distributed by disrupting the communication channels. But even if no adversary is present, one might still wish to take into account that, due to noise or another disturbance, it can happen that no key is generated. One may in this case be able to bound the probability of successfully distributing a key, and thus the ideal resource constructed is stronger than that of Fig. 6(b) (where there is no bound on the probability of getting a key) but weaker than that of Fig. 6(a) (where a key is generated with probability 1). This middle point is depicted in Fig. 6(c) (where a key is generated with probability $1-\delta$) and is treated in Sec. III.B.4.

2. Real QKD system

a. Protocol

There are various types of QKD protocols that differ by their use of resources and hence practical feasibility (Scarani *et al.*, 2009). For example, in entanglement-based protocols, which were first proposed by Ekert (1991), Alice and Bob use a source of entanglement together with a classical authentic (but otherwise insecure) communication channel to generate their keys. Here we focus on prepare-and-measure schemes, where instead of having access to entanglement it is assumed that Alice can send quantum states to Bob. These protocols, of which that of Bennett and Brassard (1984) is the most prominent example, are technologically less challenging than entanglement-based ones, for they do not require the generation of entanglement. Alice merely has to prepare states and send them to Bob, and Bob has to measure them.

QKD protocols can be roughly divided into three phases: quantum state distribution, error estimation, and classical postprocessing. In the first, Alice sends some quantum states to Bob, who measures them upon reception, obtaining a classical string, called the raw key. In the error estimation phase, they sample some bits at random positions in the raw key and estimate the noise on the quantum channel by comparing these values to what Bob should have obtained. If the noise level is above a certain threshold, they abort the protocol and output an error message. If the noise is low enough, they move on to the third phase, in which they perform error correction and privacy amplification on their respective strings. Error correction allows Bob to correct the bits where his raw key differs from Alice's. Privacy amplification turns the raw key, about which an adversary may still have partial information, into the final secret key, i.e., uniform strings k_A and k_B for Alice and Bob, respectively (which ideally should be equal).

b. Resources

The security of a QKD protocol depends also on the resources that we start with. As previously mentioned, we

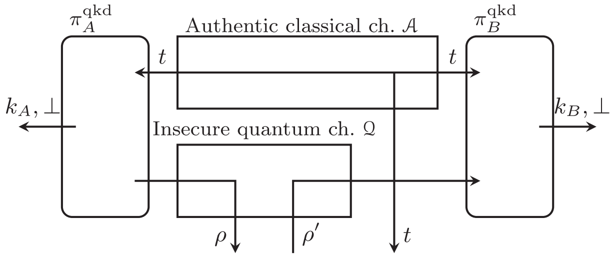

(a)

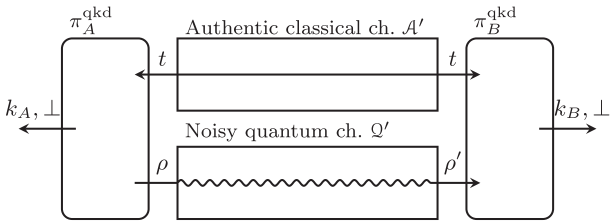

(b)

FIG. 7. The real QKD system (Alice has access to the left interface, Bob to the right interface, and Eve to the lower interface) consists of the protocol $(\pi_A^{\mathrm{QKD}}, \pi_B^{\mathrm{QKD}})$, the insecure quantum channel $\mathcal{Q}$ in (a) [a noisy channel $\mathcal{Q}'$ in (b)], and the two-way authentic classical channel $\mathcal{A}$ [$\mathcal{A}'$ in (b)]. As before, arrows represent the transmission of classical or quantum messages. The protocols of Alice and Bob $(\pi_A^{\mathrm{QKD}}, \pi_B^{\mathrm{QKD}})$ abort if they detect too much interference, i.e., if ρ' is not similar enough to ρ to obtain a secret key of the desired length. They run the classical postprocessing over the authentic channel, obtaining keys k_A and k_B. The message t depicted on the two-way authentic channel represents the entire transcript of the classical communication between Alice and Bob during the protocol. (a) Eve's interfaces of the channel resources give her full access to the quantum communication and allow her to read the messages on the authentic channel. (b) In a model with natural noise, the resources $\mathcal{Q}$ and $\mathcal{A}$ are replaced by nonmalicious variants $\mathcal{Q}'$ and $\mathcal{A}'$ that have a blank interface for Eve and a fixed noise model for the channel $\mathcal{Q}'$.

are interested in making statements about two cases. In the presence of an active adversary, we want to guarantee that any key generated is secure (soundness). But this is not sufficient, since a protocol that always aborts and never distributes a key satisfies such a criterion but is pointless. We thus also want to guarantee that if no adversary is present [only natural (low) noise] a key will be generated with high probability (completeness).

These two cases are modeled by considering different resources in the real world. In the case of an active adversary, the resources available for a prepare-and-measure scheme are a one-way insecure quantum channel from Alice to Bob (i.e., Eve may change and insert messages on the channel) and a classical two-way authentic channel (i.e., it allows authenticated communication from Alice to Bob and Bob to Alice, but Eve may also listen in). These are illustrated in Fig. 7(a). Recall that this construction is then supposed to realize the ideal system depicted in Fig. 6(b).

The quantum channel is used in the protocol when Alice sends the qubits she prepared to Bob. This channel may be

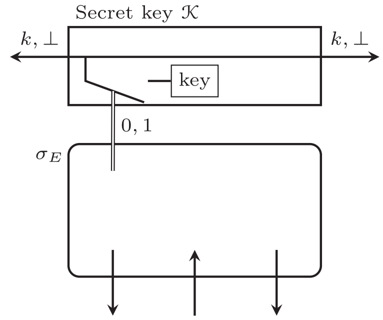

FIG. 8. Key resource from Fig. 6(b) with a simulator σ_E. This corresponds to the ideal world in Eq. (8).

completely under the control of Eve, who could apply any operation allowed by physics to what is sent over the channel. The authentic channel is used during the next two phases of the protocol, in which Alice and Bob estimate the noise in their raw keys and perform the postprocessing. Such a channel faithfully transmits messages between Alice and Bob but provides Eve with a copy as well. Since an authentic channel can be constructed from an insecure channel and a short shared secret key, QKD is sometimes referred to as a key expansion protocol.[19]

The second case is modeled by resources that are no longer controlled by Eve. Instead, the quantum channel has a fixed noise model and the authentic channel does not provide copies of the messages to Eve. This is drawn in Fig. 7(b). With these assumed resources, the ideal resource one wishes to construct is given by Fig. 6(c).

3. Security

For the following, we denote by $(\pi_A^{\mathrm{QKD}}, \pi_B^{\mathrm{QKD}})$ the QKD protocol, with π_A^{QKD} and π_B^{QKD} the converters applied by Alice and Bob, respectively. We furthermore denote by $\mathcal{Q}$ the insecure quantum channel and by $\mathcal{A}$ the authentic classical channel, as drawn in Fig. 7(a). Their nonmalicious counterparts are denoted by $\mathcal{Q}'$ and $\mathcal{A}'$, respectively, as in Fig. 7(b). Finally, let $\mathcal{K}$ be the secret key resource of Fig. 6(b), and let $\mathcal{K}'$ be the secret key resource of Fig. 6(c). Applying Definition 1, we find that $(\pi_A^{\mathrm{QKD}}, \pi_B^{\mathrm{QKD}})$ constructs $\mathcal{K}$ from $\mathcal{Q}$ and $\mathcal{A}$ within ε if

$$\exists\, \sigma_E \qquad \pi_A^{\mathrm{QKD}} \pi_B^{\mathrm{QKD}} (\mathcal{Q} \| \mathcal{A}) \approx_\varepsilon \mathcal{K}\sigma_E, \tag{8}$$

and $(\pi_A^{\mathrm{QKD}}, \pi_B^{\mathrm{QKD}})$ constructs $\mathcal{K}'$ from $\mathcal{Q}'$ and $\mathcal{A}'$ within ε' if

$$\pi_A^{\mathrm{QKD}} \pi_B^{\mathrm{QKD}} (\mathcal{Q}' \| \mathcal{A}') \approx_{\varepsilon'} \mathcal{K}'. \tag{9}$$

Note that no simulator is needed in Eq. (9) because both the real and ideal systems have a blank interface for Eve. The left- and right-hand sides of Eq. (8) are illustrated in Figs. 7(a) and 8, and the left- and right-hand sides of Eq. (9) are illustrated in Figs. 7(b) and 6(c). These two conditions are decomposed into simpler criteria in Sec. III.B.

[19] We model QKD this way in Sec. VII.

B. Reduction to the trace distance criterion

By applying the general AC security definition to QKD, we obtained two criteria, Eqs. (8) and (9), capturing soundness and completeness, respectively. In this section we derive the trace distance criterion, Eq. (7), introduced at the beginning of Sec. III, from Eq. (8). We first show in Sec. III.B.1 that the distinguishing advantage used previously in the review reduces to the trace distance between the quantum states gathered by the distinguisher interacting with the real and ideal systems. In Sec. III.B.2 we then determine the simulator σ_E of the ideal system. In Sec. III.B.3 we decompose the resulting security criterion into a combination of secrecy (the trace distance criterion) and correctness (the probability that Alice's and Bob's keys differ). In Sec. III.B.4 we consider the security condition of Eq. (9), which captures security guarantees in the absence of a malicious adversary. We show how this condition can be used to model the robustness of the protocol, i.e., the probability that the protocol aborts with nonmalicious noise.

1. Trace distance

The security criteria given in Eqs. (8) and (9) are defined in terms of the distinguishing advantage between resources. To simplify these equations, we rewrite them in terms of the trace distance between the states held by the distinguisher at the end of the protocol in the real and ideal settings. Helstrom (1976) proved that the advantage a distinguisher has in guessing whether it was provided with one of two states with equal priors, ρ or σ, is given by the trace distance between the two $D(\rho, \sigma)$.[20] A proof of this along with a discussion of different operational interpretations of the trace distance is given in Appendix A.

We start with the criterion given by Eq. (9). The two resources on the left- and right-hand sides of Eq. (9) simply output classical strings (a key or error message) at Alice's and Bob's interfaces. Let these pairs of strings be given by the joint probability distributions P_{AB} and $\tilde{P}_{AB}$. The distinguishing advantage between the two resources is thus simply the distinguishing advantage between these probability distributions (a distinguisher is given a pair of strings sampled according to either P_{AB} or $\tilde{P}_{AB}$ and has to guess from which distribution it was sampled), i.e.,

$$d(\pi_A^{\mathrm{QKD}}\pi_B^{\mathrm{QKD}}(\mathcal{Q}'\|\mathcal{A}'), \mathcal{K}') = d(P_{AB}, \tilde{P}_{AB}).$$

As previously stated, the distinguishing advantage between two quantum states is equal to their trace distance, and in the special case where the states are classical (i.e., given by two probability distributions) the trace distance between the classical states is equal to the total variational distance between the corresponding probability distributions. Thus, $d(P_{AB}, \tilde{P}_{AB}) = D(P_{AB}, \tilde{P}_{AB})$, where we use the same notation for both the trace distance and the total variational distance,

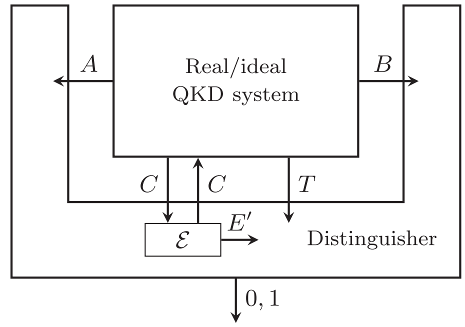

FIG. 9. The distinguisher interacting with either the real or ideal QKD system first receives a register C containing the quantum states sent from Alice to Bob. It applies a map $\mathcal{E}: \mathcal{L}(\mathcal{H}_C) \to \mathcal{L}(\mathcal{H}_{CE'})$ of its choice, keeps the E' register, and puts C back in the insecure channel. Finally, it gets the transcript of the classical communication T, and Alice's and Bob's outputs A and B. It thus holds a state $\rho_{ABE'T}$, which it measures to decide whether it was interacting with the real or ideal system.

since the latter is a special case of the former. Putting the two together we get

$$d(\pi_A^{\mathrm{QKD}}\pi_B^{\mathrm{QKD}}(\mathcal{Q}'\|\mathcal{A}'), \mathcal{K}') = D(P_{AB}, \tilde{P}_{AB}),$$

where P_{AB} and $\tilde{P}_{AB}$ are the distributions of the string output by the real and ideal systems, respectively.

The resources on the left- and right-hand sides of Eq. (8) are slightly more complex than those in Eq. (9). They first output a state φ_C at the E interface, namely, the quantum states that Alice sends over the insecure quantum channel. Without loss of generality, the distinguisher now applies any map $\mathcal{E}: \mathcal{L}(\mathcal{H}_C) \to \mathcal{L}(\mathcal{H}_{CE'})$ allowed by quantum physics to this state, obtaining $\rho_{CE'} = \mathcal{E}(\varphi_C)$, and puts the C register back on the insecure channel for Bob, keeping the part in E'. Finally, the systems output some keys (or error messages) at the A and B interfaces and all classical messages exchanged during the error estimation and postprocessing at the E interface: this captures the fact that the classical communication is public.[21] This sequence of interactions of the distinguisher with the real or ideal QKD systems is illustrated in Fig. 9.

Let $\rho_{ABE}^{\mathcal{E}}$ be the tripartite state held by a distinguisher interacting with the real system, and let $\tilde{\rho}_{ABE}^{\mathcal{E}}$ be the state held after interacting with the ideal system, where the registers A and B contain the final keys or error messages.The register E then holds both the state $\rho_{E'}$ obtained from tampering with the quantum channel and the classical transcript. Distinguishing between these two systems thus reduces to maximizing over the distinguisher strategies (the choice of $\mathcal{E}$) and distinguishing between the resulting states $\rho_{ABE}^{\mathcal{E}}$ and $\tilde{\rho}_{ABE}^{\mathcal{E}}$ as follows:

$$d(\pi_A^{\mathrm{QKD}}\pi_B^{\mathrm{QKD}}(\mathcal{Q}\|\mathcal{A}), \mathcal{K}\sigma_E) = \max_{\mathcal{E}} d(\rho_{ABE}^{\mathcal{E}}, \tilde{\rho}_{ABE}^{\mathcal{E}}).$$

Using again the equality between trace distance and distinguishing advantage, we obtain that the advantage a

[20]Actually, Helstrom (1976) solved a more general problem in which the states ρ and σ are picked with *a priori* probabilities p and $1 - p$, respectively, instead of $1/2$ as in the definition of the distinguishing advantage.

[21]We sometimes refer to the entire sequence of these messages as the classical transcript of the protocol.

distinguisher has in guessing whether it holds the state $\rho^{\mathcal{E}}_{ABE}$ or $\tilde{\rho}^{\mathcal{E}}_{ABE}$ is given by the trace distance between these states, i.e.,

$$d(\pi_A^{\mathrm{QKD}}\pi_B^{\mathrm{QKD}}(\mathcal{Q}\|\mathcal{A}), \mathcal{K}\sigma_E) = \max_{\mathcal{E}} D(\rho^{\mathcal{E}}_{ABE}, \tilde{\rho}^{\mathcal{E}}_{ABE}).$$

The distinguishing advantage between the real and ideal systems of Eq. (8) thus reduces to the trace distance between the quantum states gathered by the distinguisher. In the following, we usually omit $\mathcal{E}$ when it is clear that we are maximizing over the distinguisher strategies, and simply express the security criterion as

$$D(\rho_{ABE}, \tilde{\rho}_{ABE}) \leq \varepsilon, \tag{10}$$

where ρ_{ABE} and $\tilde{\rho}_{ABE}$ are the quantum states gathered by the distinguisher interacting with the real and ideal systems, respectively.

2. Simulator

In the real setting [Fig. 7(a)], Eve has full control over the quantum channel and obtains the entire classical transcript of the protocol. Thus, for the real and ideal settings to be indistinguishable, a simulator σ_E^{QKD} must generate the same communication as in the real setting. This can be done by internally running Alice's and Bob's protocols $(\pi_A^{\mathrm{qkd}}, \pi_B^{\mathrm{qkd}})$, producing the same messages at Eve's interface as the real system. However, instead of letting this simulated protocol decide the value of the key as in the real setting, the simulator ignores these values and checks only whether a key is actually produced or an error message is generated instead. It then operates the switch on the secret key resource accordingly. We illustrate this in Fig. 10.

The security criterion from Eq. (10) can now be simplified by noting that with this simulator the states of the ideal and real systems are identical when no key is produced. The outputs at Alice's and Bob's interfaces are classical elements of the set $\{\perp\} \cup \mathcal{K}$, where $\perp$ symbolizes an error and $\mathcal{K}$ is the set of possible keys. The states of the real and ideal systems can be written as

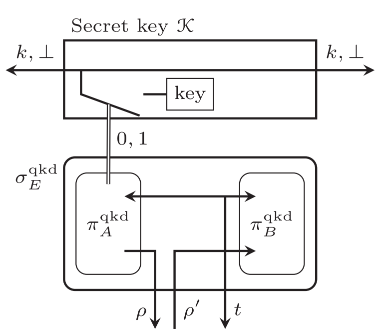

FIG. 10. The ideal QKD system (Alice has access to the left interface, Bob has access to the right interface, and Eve has access to the lower interface) consists of the ideal secret key resource and a simulator σ_E^{QKD}.

$$\begin{aligned}\rho_{ABE} &= p^{\perp}|\perp_A,\perp_B\rangle\langle\perp_A,\perp_B| \otimes \rho_E^{\perp} \\ &\quad + \sum_{k_A,k_B\in\mathcal{K}} p_{k_A,k_B}|k_A,k_B\rangle\langle k_A,k_B| \otimes \rho_E^{k_A,k_B}, \\ \tilde{\rho}_{ABE} &= p^{\perp}|\perp_A,\perp_B\rangle\langle\perp_A,\perp_B| \otimes \rho_E^{\perp} \\ &\quad + \frac{1}{|\mathcal{K}|}\sum_{k\in\mathcal{K}}|k,k\rangle\langle k,k| \otimes \sum_{k_A,k_B\in\mathcal{K}} p_{k_A,k_B}\rho_E^{k_A,k_B},\end{aligned}$$

where p_{k_A,k_B} is the probability of Alice getting the key k_A and Bob getting k_B and $p^{\perp}$ is the probability of an abort. Plugging this into Eq. (10), we get

$$D(\rho_{ABE}, \tilde{\rho}_{ABE}) = (1-p^{\perp})D(\rho^{\top}_{ABE}, \tau_{AB}\otimes\rho^{\top}_E) \leq \varepsilon, \tag{11}$$

where

$$\rho^{\top}_{ABE} \coloneqq \frac{1}{1-p^{\perp}} \sum_{k_A,k_B\in\mathcal{K}} p_{k_A,k_B}|k_A,k_B\rangle\langle k_A,k_B| \otimes \rho_E^{k_A,k_B} \tag{12}$$

is the renormalized state of the system conditioned on not aborting and $\tau_{AB} \coloneqq (1/|\mathcal{K}|)\sum_{k\in\mathcal{K}}|k,k\rangle\langle k,k|$ is a perfectly uniform shared key. As previously, the E register contains the quantum side information that Eve collects about the states being sent as well as the entire classical transcript of the error estimation and postprocessing.

3. Correctness and secrecy

We now break up Eq. (11) into two components, often referred to as correctness and secrecy, and recover the security definition for QKD that was introduced by Ben-Or *et al.* (2005), Renner (2005), and Renner and König (2005). The correctness of a QKD protocol refers to the probability that Alice and Bob end up holding different keys. We say that a protocol is $\varepsilon_{\mathrm{cor}}$ correct if for all adversarial strategies

$$\Pr[K_A \neq K_B] \leq \varepsilon_{\mathrm{cor}}, \tag{13}$$

where K_A and K_B are random variables over the alphabet $\mathcal{K}\cup\{\perp\}$ describing Alice's and Bob's outputs.[22] The secrecy of a QKD protocol measures how close the final key is to a distribution that is uniform and independent of the adversary's system. Let $p^{\perp}$ be the probability that the protocol aborts, and let $\rho^{\top}_{AE}$ be the resulting state of the AE subsystems conditioned on not aborting. A protocol is $\varepsilon_{\mathrm{sec}}$ secret if for all adversarial strategies

$$(1-p^{\perp})D(\rho^{\top}_{AE}, \tau_A\otimes\rho^{\top}_E) \leq \varepsilon_{\mathrm{sec}}, \tag{14}$$

where the distance $D(\cdot,\cdot)$ is the trace distance and τ_A is the fully mixed state.[23]

[22]This can equivalently be written as $(1-p^{\perp})\Pr[K_A^{\top} \neq K_B^{\top}] \leq \varepsilon_{\mathrm{cor}}$, where $p^{\perp}$ is the probability of aborting and $K_A^{\top}$ and $K_B^{\top}$ are Alice's and Bob's keys conditioned on not aborting.

[23]Equation (14) was already introduced at the beginning of Sec. III as Eq. (7).

Theorem 2.—If a QKD protocol is $\varepsilon_{\mathrm{cor}}$ correct and $\varepsilon_{\mathrm{sec}}$ secret, then Eq. (8) is satisfied for $\varepsilon = \varepsilon_{\mathrm{cor}} + \varepsilon_{\mathrm{sec}}$.

This theorem can be proven by using the triangle inequality of the trace distance to bound Eq. (11) in terms of the sum of correctness and secrecy. For completeness, a proof is given in Appendix B. This result was also given by Ben-Or *et al.* (2005).

The converse statement can also be shown: if Eq. (8) holds for some ε, then the corresponding QKD protocol is both ε correct and 2ε secret.[24]

4. Robustness

Correctness and secrecy, as previously described, capture the soundness of QKD in the presence of a malicious Eve, as specified by Eq. (8). This, however, is not sufficient: a QKD protocol that always aborts without producing any key trivially satisfies Eq. (8) with $\varepsilon = 0$ but is not at all a useful protocol. This is where the second condition, namely, Eq. (9), is relevant. The real system must be indistinguishable from ideal not only when an adversary is present and manipulating the channel but also when one has a simple noisy channel, with a blank adversarial interface. In this case, we expect a secret key to be generated successfully with high probability. This is captured by considering the strong ideal key resource $\mathcal{K}'$ from Fig. 6(c), which produces a key with probability $1-\delta$. If the real system does not generate a key with the same probability, this immediately results in a gap that is noticeable to the distinguisher.

The probability that the real protocol generates a key depends on the noise introduced by the noisy channel $\mathcal{Q}'$ [illustrated in Fig. 7(b)]. Suppose that this noise is parametrized by a value q such as a depolarizing channel with probability q. For every q, the protocol has a probability of aborting δ, which is called the robustness. Let $\mathcal{Q}_q$ denote a channel with this noise model, and let $\mathcal{K}_\delta$ denote the key resource that produces an error with a fixed probability δ. Equation (9) can thus be phrased as

$$\pi_A^{\mathrm{QKD}}\pi_B^{\mathrm{QKD}}(\mathcal{Q}_q\|\mathcal{A}') \approx_\varepsilon \mathcal{K}_\delta, \tag{15}$$

where varying q and δ results in a family of real and ideal systems.

One can then show that the failure ε from Eq. (15) is bounded by $\varepsilon_{\mathrm{cor}} + \varepsilon_{\mathrm{sec}}$. Note that this statement is useful only if the probability of aborting δ is small for reasonably noisy models q.

Lemma 3.—If the resources from Eq. (15) are parametrized such that $\mathcal{K}_\delta$ aborts with exactly the same probability as the protocol $(\pi_A^{\mathrm{QKD}}, \pi_B^{\mathrm{QKD}})$ that is run on the noisy channel $\mathcal{Q}_q$, then the completeness of the protocol is bounded by the soundness, i.e.,

$$d(\pi_A^{\mathrm{QKD}}\pi_B^{\mathrm{QKD}}(\mathcal{Q}_q\|\mathcal{A}'), \mathcal{K}_\delta) \le d(\pi_A^{\mathrm{QKD}}\pi_B^{\mathrm{QKD}}(\mathcal{Q}\|\mathcal{A}), \mathcal{K}\sigma_E^{\mathrm{QKD}}),$$

where the simulator σ_E^{QKD} is the one used previously in the review and introduced in Sec. III.B.2, Fig. 10.

A proof of this lemma is provided in Appendix B.

C. Other security criteria

1. Accessible information

As mentioned at the beginning of this section, the trace distance criterion was introduced only in 2005 (Ben-Or *et al.*, 2005; Renner, 2005; Renner and König, 2005). Earlier works (Mayers, 1996); Biham *et al.*, 2000; Shor and Preskill, 2000) used a notion of security directly inspired from classical cryptography, where key techniques such as advantage distillation, error correction, and privacy amplification were developed (Bennett, Brassard, and Robert, 1988; Ahlswede and Csiszár, 1993; Maurer, 1993; Bennett *et al.*, 1995). More concretely, if one denotes an n-bit key random variable by K and denotes the adversary's classical side information by Z, in these works a key was considered secure if the mutual information per bit between the two is small, i.e., $(1/n)I(K;Z) \le \varepsilon$, where $I(K;Z) = H(K) - H(K|Z)$. It was later realized (Maurer, 1994; Maurer and Wolf, 2000) that the mutual information per bit is not appropriate in the asymptotic setting, since $\varepsilon(n) \to 0$ does not imply that the total information about the key is also small; i.e., one may still have $n\varepsilon(n) \nrightarrow 0$. It was therefore considered preferable to directly bound the total information about the key $[I(K;Z) \le \varepsilon]$.

In the case of QKD the side information may be quantum, and the joint system of the key and side information is given by a state ρ_{KE}. The accessible information between K and E is obtained by measuring the E system and taking the mutual information between K and the measurement outcome, i.e.,

$$I_{\mathrm{acc}}(K;E)_\rho \coloneqq \max_{\{\Gamma_z\}_z} I[K;\Gamma_Z(E)] \le \varepsilon, \tag{16}$$

where $\Gamma_Z(E)$ is the random variable resulting from a measurement of the E system with the positive operator-valued measure (POVM) $\{\Gamma_z\}_z$ and, as before, $I(K;Z) = H(K) - H(K|Z)$ is the mutual information.

Since measuring a quantum system can only diminish the information that it provides, one always has $S(K|E) \le H(K|Z)$ for any random variable Z obtained by measuring the E system of a bipartite state ρ_{KE}, where $S(\cdot)$ is the von Neumann entropy. Using the continuity of the conditional von Neumann entropy (Alicki and Fannes, 2004), this can by bounded by its trace distance from uniform, namely,[25]

$$n - S(K|E) \le 8\delta n + 2h(2\delta),$$

where δ is the trace distance between ρ_{KE} and $\tau_K \otimes \rho_E$ and $h(p) = -p\log p - (1-p)\log(1-p)$ is the binary entropy.

[24] The factor 2 is due to the existence quantifier over simulators σ_E in the security definition. We cannot exclude that for some specific QKD protocol there is a simulator $\bar{\sigma}_E^{\mathrm{QKD}}$, unlike the one used in this proof, that generates a state $\bar{\rho}_E$ satisfying $D(\rho_{AE}^\top, \tau_A \otimes \bar{\rho}_E^\top) \le D(\rho_{AE}^\top, \tau_A \otimes \rho_E^\top)$. However, by the triangle inequality we also find that, for any $\bar{\rho}_E$, $D(\rho_{AE}^\top, \tau_A \otimes \bar{\rho}_E^\top) \ge (1/2)D(\rho_{AE}^\top, \tau_A \otimes \rho_E^\top)$. Hence, the failure ε of the generic simulator used in this proof cannot be more than twice as large than the optimal one.

[25] See Corollary 13 in Appendix A for a proof of this.

The trace distance criterion thus provides a bound on the accessible information.

However, the converse does not hold. As König *et al.* (2007) showed, it is possible to find a joint state ρ_{KE} of an n-bit key K and the adversary's information E that satisfies Eq. (16) with $\varepsilon = 2^{-0.18n}$, but knowledge of the first $n-1$ bits K_1 of $K = K_1 K_2$ allow the last bit K_2 to be guessed perfectly.[26] More precisely, if one knows K_1, there is a way to measure the quantum system E such that the outcome $\Gamma_{Z'}(K_1 E)$ is a perfect guess[27] for K_2, i.e.,

$$I[K_2; \Gamma_{Z'}(K_1 E)] = 1.$$

To see why this is problematic, suppose that the two parts of the key K_1 and K_2 are used for one-time pad encryption (see Sec. II.C) of two messages M_1 and M_2, respectively, of which the first is already known to an eavesdropper (because it contains some publicly available information). Given M_1, the eavesdropper can, by listening to the ciphertext, infer K_1. This, in turn, allows her to apply the measurement yielding Z' to E, which provides information about K_2, and hence also about M_2.

This example emphasizes the relevance of composability, i.e., the principle that any reasonable notion of security should have the property that if two cryptographic schemes are considered secure then this should also be the case for their combination. Criterion (16) does not satisfy this principle. If a key K generated by a QKD protocol satisfies Eq. (16), then by definition it is guaranteed that an adversary cannot infer K. But the composition of this QKD protocol with the one-time pad encryption, which by itself is a perfectly secure protocol, is insecure. This is problematic, for such compositions of protocols are ubiquitous in cryptography.

To see how the real-world ideal-world paradigm avoids this issue, imagine a protocol that generates a key consisting of two parts K_1 and K_2, with the undesirable properties as in the example used by König *et al.* (2007) as previously described. The distinguisher could then use K_1 to measure E and check to see whether the outcome Z' determines K_2. If this is the case, then the distinguisher knows that it was interacting with the real system ($B = 0$); otherwise, it must have been the ideal one ($B = 1$). The distinguisher could thus correctly guess the bit B, i.e., the protocol would not meet the criterion of being indistinguishable from an ideal system. Hence, although the key generation protocol may still satisfy earlier criteria such as Eq. (16), it would be considered insecure, as it should be.

It is interesting nonetheless to understand which construction a security definition like the accessible information corresponds to. We discuss this in Sec. VI.F, where we show that if one assumes that the adversary has no quantum memory, then the accessible information is a sufficient security criterion.

[26]This phenomenon is known as *information locking*, and further examples were given by DiVincenzo *et al.* (2004) and Winter (2017).

[27]Ben-Or *et al.* (2005) showed that if $\varepsilon < 2^{-n}$ then information locking cannot be exploited, and the adversary's advantage in guessing K_2 remains exponentially small.

2. Adversarial models

The definition of cryptographic security introduced in Sec. II, Definition 1 does not explicitly mention an adversary. The notion of an adversary is embedded in the distinguisher, which is used to measure the distance between real and ideal systems. The distinguishing metric thus has a dual role: performing the most powerful attack possible and measuring whether this attack was successful, i.e., whether it allows real and ideal systems to be distinguished. The reduction to the trace distance criterion discussed in Sec. III.B separates these two notions. The trace distance criterion itself [Eq. (7)] can be seen as the measure of whether the attack resulting in the adversary holding the system E is successful. For this condition to make sense as a security definition, one has to consider all possible adversarial behaviors, i.e., take the maximum of Eq. (7) over all possible states ρ_{KE} that may occur.

Historically, these two aspects (the attack and the criterion for measuring whether the attack is successful) were treated separately. Early security proofs for QKD, such as that of Bennett, Bessette *et al.* (1992), did not consider the most powerful attack an eavesdropper could perform, but instead only individual attacks. These are attack strategies where the adversary performs an identical operation on each qubit on the quantum channel and keeps only classical information Z. The information held by Alice, Bob, and Eve is then modeled by independent and identically distributed (i.i.d.) random variables.

Collective attacks, a generalization of individual attacks that allows the eavesdropper to keep quantum information but still forces her to perform the same operation on every qubit, were proposed by Biham and Mor (1997) and Biham *et al.* (2002). In this setting, one has to use von Neumann entropy instead of Shannon entropy to measure the adversary's information about the raw key and compute the achievable rate (Devetak and Winter, 2005). Although the adversary's interactions with the quantum channel are restricted to i.i.d. operations, this class of attacks is particularly important since proof techniques developed later (Renner, 2005, 2007; Christandl, König, and Renner, 2009; Arnon-Friedman, Renner, and Vidick, 2019; Dupuis, Fawzi, and Renner, 2020) show how one can reduce the most general attack strategies to such a limited one.

The first security proof for QKD that considered a fully general adversary (performing coherent attacks) was attributed to Mayers (1996, 2001). It was then followed by other, simpler proofs (Biham *et al.*, 2000, 2006; Shor and Preskill, 2000). Biham *et al.* (2000, 2006) and Shor and Preskill (2000) pointed out that security does not hold that is conditional on the protocol terminating with a secret key. Instead, one should prove that the probability of the event that the protocol does not abort and that the adversary has nontrivial information about the key is negligible. However, the previously discussed works still use basically classical security definitions, such as those based on the accessible information.

3. Expressing weaker security criteria within the AC framework

As previously discussed, the early security definitions implicitly imposed a restriction on the set of possible attack strategies that an adversary could pursue. Within the modern real-world ideal-world paradigm, or more precisely the AC framework, one can understand these restrictions as limitations on the distinguisher that is trying to guess whether it is interacting with the real or ideal system. That is, one does not consider the full set of possible distinguishers, but instead only a restricted subset that performs i.i.d. operations or takes its final decision by measuring the E system alone, not the joint KE system.

Alternatively, one may also represent these definitions in the AC framework by either replacing the resources in the ideal setting with weaker ones or replacing the resources available in the real setting with stronger ones. To illustrate the latter, recall that in the description in Sec. III.A the insecure quantum channel used by Alice and Bob allows Eve to perform arbitrary operations on the quantum messages sent. If instead one would provide the players with a stronger resource that allows Eve to perform only i.i.d. operations or allows her to access classical information only, one would recover weaker security definitions. This is developed in more detail in Sec. VI.F.

Using such an approach, one may still regard the older security definitions as "composable," provided one is aware of the fact that the real resources are now weaker. In other words, the weaker definitions do not guarantee that a secret key is obtained from an authentic classical channel and a completely insecure quantum channel but still ensure that a secret key is obtained from a certain less insecure quantum channel that limits Eve's tampering.

We note that this approach is applicable much more generally, i.e., beyond quantum cryptography. For example, the notion of security known as *stand-alone* (Goldreich, 2004) makes the assumption that the dishonest party does not interact with the environment during the execution of the protocol. By introducing a resource that restricts the distinguisher's behavior accordingly, one can show this security definition to actually guarantee security, albeit only in a setting where honest parties have access to such a resource. Similarly, early definitions of blindness in delegated quantum computing (DQC) (Broadbent, Fitzsimons, and Kashefi, 2009; Fitzsimons and Kashefi, 2017) are not known to construct the expected ideal resource for DQC: one that takes the input from the client, only leaks a bound on the computation size to the server, and returns a possibly wrong computation result to the client (Dunjko *et al.*, 2014). However, they do construct a weaker resource that does not provide the honest player with the result of the computation. This example is discussed again in Sec. VIII.C. A further example are results in the bounded storage model by Unruh (2011), which obtains composable security if one limits the number of times a protocol is run; we discuss these further in Sec. IX.C.

4. Asymptotic versus finite-size security

The trace distance criterion [Eq. (7)] was introduced by Ben-Or *et al.* (2005), Renner (2005), and Renner and König (2005), and the relation to composable security frameworks was discussed by Ben-Or *et al.* (2005) and Müller-Quade and Renner (2009); see also Appendix A. Security proofs for QKD with respect to this criterion were developed at the same time (Christandl, Renner, and Ekert, 2004; Renner, 2005; Renner, Gisin, and Kraus, 2005). Although these newer security proofs arguably use the right security definition, they prove security only asymptotically. This means that instead of computing the failure ε for specific parameters of the protocol, one shows that $\varepsilon(n) \to 0$ when $n \to \infty$, where n is a parameter that quantifies a certain resource, typically the number of quantum signals sent during the protocol. This does not allow the failure to be evaluated for any implementation of the protocol, since implementations must necessarily generate a key in finite time, and hence with finite resources.

In the asymptotic setting one often demands that the function $\varepsilon(n)$ be negligible, i.e., smaller than $1/p(n)$ for any polynomial $p(\cdot)$; for example, $\varepsilon(n)$ could be exponentially small in n. The reasoning is usually that honest players are polynomially bounded, so they will never run the protocol more than $p(n)$ times and the accumulated error $p(n)\varepsilon(n)$ is then still negligible. Although such a requirement is standard in cryptography, it is not directly useful for practical purposes, as previously indicated. For example, a protocol with a failure given by a function $\varepsilon(n)$ that is exponentially small for $n \geq 10^{10^{10}}$ but equal to 1 otherwise, where n is the number of signals exchanged between the players, is asymptotically secure, and yet completely insecure for any realistic parameters. Conversely, the function $\varepsilon(n) = 10^{-18}$ is considered insecure in the asymptotic setting, but it guarantees that the protocol can be run once per second for the lifetime of the Universe and still have an accumulated error substantially smaller than 1.

This illustrates that asymptotic security claims can be highly ambiguous. It is thus necessary to prove finite security bounds if one wants to actually use a cryptographic scheme. This was done for basic protocols by Inamori, Lütkenhaus, and Mayers (2007), Scarani and Renner (2008), Sheridan, Thinh, and Scarani (2010), Hayashi and Tsurumaru (2012), Tomamichel *et al.* (2012), and Tomamichel and Leverrier (2017). For more advanced protocols, which are specifically designed to be implementable with imperfect hardware, finite-size security claims were given by Lim *et al.* (2014) for decoy-state QKD (which is discussed in Sec. IV.D), by Yin and Chen (2019) for twin-field QKD (Lucamarini *et al.*, 2018), and by Curty *et al.* (2014) for measurement device–independent QKD (which is discussed in Sec. VI.E).

5. Variations of the trace distance criterion

The following alternative definition for ε secrecy was proposed in the literature in place of the trace distance criterion (Tomamichel *et al.*, 2010, 2012):

$$(1 - p_{\text{abort}}) \min_{\sigma_E} D(\rho_{KE}, \tau_K \otimes \sigma_E) \leq \varepsilon. \tag{17}$$

This alternative notion is equivalent to the standard definition of secrecy [Eq. (7)] up to a factor of 2; hence, any QKD scheme proven secure with one definition is still secure according to the other, with a minor adjustment of the failure parameter ε. However, we do not know how to derive this alternative notion from a composable framework. In particular, it is not clear whether the failure ε from Eq. (17) is additive under parallel composition. For example, the concatenation of two keys that each individually satisfy Eq. (17) could possibly have a distance[28] from uniform greater than 2ε.

IV. ASSUMPTIONS FOR SECURITY

The security of a quantum cryptographic protocol relies on assumptions about the physics of the devices that are employed to implement the protocol. In this section, we discuss these assumptions. For concreteness, we focus on the case of QKD, for which we describe the full set of assumptions in Sec. IV.A. We then explain why these assumptions are needed and to what extent they are justified in Sec. IV.B. Experimental work in QKD has shown, however, that the assumptions are often difficult to meet, and actually are not met in many cases. This fact can be exploited by quantum hacking attacks, which are described in Sec. IV.C. Finally, in Sec. IV.D we discuss countermeasures against these attacks.

A. Standard assumptions for QKD

The security of QKD protocols usually relies on the following assumptions.

(1) All devices used by Alice and Bob, as well as the communication channels connecting them, are correctly and completely[29] described by quantum theory.
(2) The channel that Alice and Bob use to exchange classical messages is authentic, i.e., it is impossible for an adversary to modify messages or insert new ones.
(3) The devices that Alice and Bob use locally to execute the steps of the protocol, such as for preparing and measuring quantum systems, do exactly what they are instructed to do.

As previously indicated, owing to the lack of proof techniques, additional assumptions were introduced in the past. A prominent example is the i.i.d. assumption, which demands that the quantum channel connecting Alice and Bob be described by a sequence of identical and independently distributed maps. Physically, this means that an adversary's interception strategy is such that each signal sent from Alice to Bob is modified in the same manner and independently of the other signals. Security under the i.i.d. assumption is called security against collective attacks; see Biham and Mor (1997), Biham *et al.* (2002), and Sec. III.C.2. Another assumption, which usually supplements the i.i.d. assumption, is that Eve stores only classical data, which she obtains by individually measuring the pieces of information she has gained from each signal sent from Alice to Bob. Since it is difficult to argue why an adversary should be restricted in that particular way, the corresponding security guarantee is rather weak. It is usually referred to as security against *individual attacks*; see Fuchs *et al.* (1997), Lütkenhaus (2000), and Sec. III.C.2.

Most modern security proofs, however, do not require such additional assumptions; i.e., they are based entirely on the previously mentione assumptions (1)–(3). This means, in particular, that the quantum channel connecting Alice and Bob can be arbitrary and may even be entirely controlled by Eve. In this case, one talks about security against general attacks, coherent attacks, or joint attacks. Sometimes the term unconditional security has appeared in the literature (Scarani *et al.*, 2009), but it is important to keep in mind that the previously listed assumptions are still necessary.

B. Necessity and justification of assumptions

Assumption (1) is often implicit, for it is a prerequisite to even describe the cryptographic scheme. It justifies the use of the formalism of quantum theory to model the different systems, such as the communication channel, including any possible attacks on them. The assumption thus captures the main idea behind quantum cryptography, namely, that an adversary is limited by the laws of quantum theory. The other two assumptions ensure that the experimental implementation follows the theoretical prescription that enters the security definition (Definition 1), namely, the description of the protocol π_{AB} and the used resources. In particular, assumption (2) guarantees that the resources shared between Alice and Bob fulfill the theoretical specifications $\mathcal{R}$, which in the case of QKD includes the classical authentic communication channel. Assumption (3) guarantees that the steps prescribed by the protocol π_{AB} are correctly executed.

Assumption (1) is widely accepted, and proving it wrong would represent a major breakthrough in physics. Nevertheless, it has been shown that there are QKD protocols that rely only on the weaker assumption of no signaling (Barrett, Hardy, and Kent, 2005).

Assumption (2) demands that an authentic communication channel is set up between Alice and Bob. There are information-theoretically secure protocols that achieve this, provided that Alice and Bob share a weak secret key; see Renner and Wolf (2003), Dodis and Wichs (2009), Aggarwal *et al.* (2019), and Sec. VII.A. Assumption (2) can thus be met with the use of such authentication protocols; see also Sec. I as well as standard textbooks on classical cryptography.

Although assumption (3) sounds rather natural and is in fact required for almost any cryptographic scheme, including any classical one, it is challenging to meet. The numerous quantum hacking experiments that have been conducted over the past few years have shown that many implementations of QKD failed to satisfy this assumption. To illustrate this problem, we describe selected examples of such attacks in Sec. IV.C.

C. Quantum hacking attacks

We start with the *photon number splitting attack* (Brassard *et al.*, 2000), which targets optical implementations of QKD that use individual photons as quantum information carriers. Suppose for concreteness that Alice and Bob implement the BB84 protocol (Bennett and Brassard, 1984) by encoding the qubits into the polarization degree of freedom of individual photons. Specifically, Alice may use a single-photon source that emits photons with a polarization that she can choose.

[28]The preprint version of Tomamichel *et al.* (2012) was updated to use Eq. (7) instead.

[29]The completeness of quantum theory can be derived from their correctness; see Sec. I.A.

The BB84 protocol[30] requires her to send in each round at random a state from one orthonormal basis, say, $\{|h\rangle, |v\rangle\}$, where $|h\rangle$ may be realized with a horizontally polarized photon and $|v\rangle$ may be realized with a vertically polarized one, or from a complementary basis $\{|d^+\rangle, |d^-\rangle\}$, where $|d^+\rangle = (1/\sqrt{2})(|h\rangle + |v\rangle)$ and $|d^-\rangle = (1/\sqrt{2})(|h\rangle - |v\rangle)$. It may now happen that, in an experimental implementation, the source sometimes accidentally emits two photons at once, which then carry the same polarization. The states emitted in the four cases are thus $|h\rangle \otimes |h\rangle$, $|v\rangle \otimes |v\rangle$, $|d^+\rangle \otimes |d^+\rangle$, and $|d^-\rangle \otimes |d^-\rangle$.

Before describing the actual attack, we first give a simple information-theoretic argument for why this is problematic. Note first that one single photon carries no information about the choice of the basis made by Alice. Indeed, for either of the basis choices, the density operator describing the photon is maximally mixed, i.e., $1/2|h\rangle\langle h| + 1/2|v\rangle\langle v| = 1/2|d^+\rangle\langle d^+| + 1/2|d^-\rangle\langle d^-| = (1/2)\mathbf{1}$. This is, however, no longer the case for a pulse consisting of two photons, i.e.,

$$\tfrac{1}{2}|h\rangle\langle h|^{\otimes 2} + \tfrac{1}{2}|v\rangle\langle v|^{\otimes 2} \neq \tfrac{1}{2}|d^+\rangle\langle d^+|^{\otimes 2} + \tfrac{1}{2}|d^-\rangle\langle d^-|^{\otimes 2}. \tag{18}$$

Hence, if the source accidentally emits two equally polarized photons instead of one, it reveals information about Alice's basis choice, which it should not.

It is therefore not surprising that such two-photon pulses can be exploited by an adversary to attack the system. Eve, who intercepts the channel, may split the two-photon pulse into two, keep one of the photons, and send the other one to Bob. The latter thus receives photons in exactly the way prescribed by the protocol, and hence does not notice the interception. Eve, meanwhile, may measure the photons she captured. In principle, if Eve had quantum memory, she could even wait with the measurement until Alice announces the basis choice to Bob, and hence could always gain full information about the polarization state that Alice prepared.

While the photon number–splitting attack exploits an imperfection of the sender (namely, that it sometimes emits two identically polarized photons instead of one), many quantum attacks are targeted toward the receiver. An example is the time-shift attack (Makarov, Anisimov, and Skaar, 2006; Qi *et al.*, 2007; Zhao *et al.*, 2008), which exploits inaccuracies of the photon detectors. To avoid dark counts, the photon detectors are often set up such that they count only those photons that arrive within a small time window around the time when a signal is expected to arrive. Furthermore, Bob's receiver device may consist of more than one detector, e.g., one for each possible polarization state. The time windows of the different detectors are then never perfectly synchronized. This means that there are times at which the receiver is more sensitive to signals with respect to one polarization than another. Eve may therefore, by appropriately delaying the signals sent from Alice and Bob, bias the detected signals toward one or the other polarization, and thus gain information about what Bob measures. While this information may be partial, it can, together with the error correction information that is available to Eve, be sufficient to infer the final key.

Another attack that is targeted toward the receiver is the *detector blinding attack* (Makarov, 2009; Lydersen *et al.*, 2010; Gerhardt *et al.*, 2011; Weier *et al.*, 2011), where the adversary tries to control the detectors by illuminating them with bright laser light. In a QKD implementation that uses the encoding of information into the polarization of individual photons, the detectors are usually configured such they can optimally detect single-photon pulses. That is, they should click whenever the incoming pulse contains a photon, and not click if the pulse is empty. However, the behavior of such detectors may be rather different in a regime where the incoming pulses contain many photons. For example, it could be that they always click when they are exposed to bright light with a particular intensity, and they may never click for another intensity. Hence, by sending in light with appropriately chosen polarization and intensity, Eve may gain immediate control over the clicks of Bob's detector. To exploit this for an attack, Eve may mimic Bob's receiver, i.e., intercept the photons sent from Alice and measure them in a randomly chosen basis, as Bob would do. She then sends bright light to Bob to ensure that he obtain the same detector clicks he would have had he directly obtained Alice's photons. This works particularly well for implementations that use a passive basis choice; i.e., where Bob's measurement basis is not provided as an input but rather made by the detection device itself. In this case, an adversary can essentially remote control Bob and thus get hold of the entire key.

Yet another hacking strategy involves Trojan-horse attacks (Vakhitov, Makarov, and Hjelme, 2001; Gisin *et al.*, 2006). Here the idea is to send a bright laser pulse via the optical fiber into Alice's or Bob's component to extract information about its internal settings. Depending on the sender and receiver hardware that is used, measuring the reflection of the pulse can allow Eve, for instance, to determine the basis choices made by Alice and Bob.

In some optical implementations of QKD, such as the plug-and-play-type (Muller *et al.*, 1997) or the circular-type (Nishioka *et al.*, 2002) system, Alice does not have a photon source but instead encodes information by modulating an incoming signal from Bob before sending it back to him. The signal thus travels twice in opposite directions through the same optical links, which helps reducing fluctuations due to birefringence and environmental noise. The twofold use of the insecure channel, however, opens additional possibilities of attacks (Gisin *et al.*, 2006). A prominent example is the phase-remapping attack (Fung *et al.*, 2007; Xu, Qi, and Lo, 2010). It exploits the fact that the modulator used by Alice to encode information into the signal coming from Bob acts on that signal during a particular time interval. In the attack, the adversary slightly advances or delays the signal on its way from Bob to Alice so that it no longer lies fully within that time interval. The modulation by Alice will then be incomplete, which means that the encoding of the information in the signal differs from what is foreseen by the protocol. This can in turn be exploited by Eve in an intercept-and-resend attack on the signal returned from Alice to Bob.

[30] This protocol is explained in more detail in Sec. V, where a security proof is also sketched.

D. Countermeasures against quantum hacking

The attacks described here have in common that they all exploit a breakdown of assumption (3). Specifically, in the case of the photon-splitting attack, the device used by Alice sends out more information than it is supposed to. In the case of the time-shift attack, it is Bob's measurement device whose measurement operators are not constant over time and can even be partially controlled by Eve. Finally, in the case of the detector blinding attack on systems with passive basis choice, Eve even takes over control of the randomness used to choose the basis.

A seemingly obvious countermeasure to prevent such attacks is to manufacture sources and detectors that meet the theoretical specifications. That is, one would need a perfect single-photon source, as well as detectors that are perfectly efficient and measure photon pulses in only a specified parameter regime. Such requirements are, however, unrealistic: the devices used in experiments will always, at least slightly, deviate from these specifications.

The other possibility is to develop cryptographic protocols and security proofs that tolerate imperfections of the devices (Gottesman *et al.*, 2004). This has been done, in particular, for the previously described attacks. To prevent photon number–splitting attacks, an efficient countermeasure is the decoy-state method (Hwang, 2003; Lo, Ma, and Chen, 2005; Wang, 2005). The idea here is that Alice sometimes deliberately sends multiphoton pulses. Alice and Bob can then check statistically to see whether an adversary captured them. Another possibility is to use protocols where Alice's encoding of information has the property that, even when one photon is extracted from a pulse, the information about what Alice sent is still partial (Scarani *et al.*, 2004; Tamaki and Lo, 2006; Sasaki, Yamamoto, and Koashi, 2014). In the case of time-shift attacks, it is sufficient to characterize the maximum bias in the detector efficiencies that can be introduced and account for it in the security proofs. Finally, for the detector blinding attacks, a possible countermeasure is to add tests to the protocol, such as a monitoring of the photocurrent, in order to detect those (Yuan, Dynes, and Shields, 2010).

The main problem with such countermeasures, however, is that the space of possible imperfections is hard to characterize. These are merely a few examples of attacks. Many others have been proposed, and have sometimes even been demonstrated to work successfully in experiments. For example, an adversary may exploit imperfections in the randomness that Alice and Bob use for choosing their measurement basis. To prevent such attacks, one may again extend the protocols such that they can tolerate imperfect randomness; see Sec. VI.C.

The last decade has thus seen an arms race between designers and attackers of quantum cryptographic schemes. A possible way out of this unsatisfactory situation is device-independent cryptography. Here the idea is to replace assumption (3) by something much weaker. Namely, one requires that the devices used by Alice and Bob do not unintentionally send out information to an adversary, and that the classical processing of information done by Alice and Bob is correct. However, one no longer demands that the sources and detectors used by Alice and Bob work according to their specifications. The way this can work is explained in Sec. VI.D.

V. SECURITY PROOFS FOR QKD

In this section, we discuss security proofs for QKD. For this we consider a generic protocol as shown in Fig. 11. The techniques presented here, however, are not restricted to QKD. Concepts such as information reconciliation or privacy amplification, which we describe in this section, also play a role in other protocols, for instance, those discussed in Sec. VIII.

While the first QKD security proofs, such as Shor and Preskill (2000) and Mayers (2001), treat the entire QKD protocol as a whole, modern security proofs are modular (Renner, 2005). This means that a separate security statement is established for each part of the cryptographic protocol. The overall security statement for QKD then follows by combining these individual statements. In the case of the protocol shown in Fig. 11, one statement concerns the raw key distribution and parameter estimation step (see Sec. V.B), another one concerns the information reconciliation step (see Sec. V.C), and yet another one concerns the privacy amplification step (see Sec. V.D). According to the AC framework, each part can be regarded as a constructive statement, which asserts that the corresponding subprotocol constructs a particular resource from certain given resources. This modular analysis comes not only with the advantage that the proofs are more versatile and can be adapted to different protocols but also that the arguments are more transparent and easier to understand and verify.

In the following, we focus on the modular approach to proving security proposed by Renner (2005). We note, however, that there are various other methods (we discuss these in Sec. V.E). The common feature of all security proofs is that they derive a relation between the information accessible to the legitimate parties and the maximum information that may have been gained by Eve. In the following description, this relation is given by Eq. (22); it lower bounds Eve's uncertainty about the raw key $\mathbf{X}$ generated by Alice. Although the statement concerns Eve's knowledge, the bound depends only on data that are accessible to Alice and Bob, in this case the error rate η_0 between their raw keys $\mathbf{X}$ and $\mathbf{Y}$.

There are various different ways to derive and interpret such bounds on Eve's information. In the case of prepare-and-measure schemes, the bounds can be understood as consequences of the no-cloning principle (Wootters and Zurek, 1982). According to this principle, if Eve attempts to copy parts of the information transmitted from Alice to Bob into her register E, the transmitted information is disturbed, resulting in a decrease of the correlations between Alice and Bob.

1. **protocol** QKD
2. $(\mathbf{X},\mathbf{Y}) := \mathrm{RawKeyDistribution}()$
3. **if** $\mathrm{ParameterEstimation}(\mathbf{X},\mathbf{Y}) = \mathrm{fail}$
   then return $(\perp,\perp)$ and **abort**
4. $(\mathbf{X},\mathbf{X}') := \mathrm{InformationReconciliation}(\mathbf{X},\mathbf{Y})$
5. $(\mathbf{S},\mathbf{S}') := \mathrm{PrivacyAmplification}(\mathbf{X},\mathbf{X}')$
6. **return** $(\mathbf{S},\mathbf{S}')$

FIG. 11. Generic QKD protocol.

protocol RawKeyDistribution() [BB84]
parameters n [number of signals]; $\phi_{x,0} := |x\rangle$, $\phi_{x,1} := \frac{1}{\sqrt{2}}(|0\rangle + (-1)^x|1\rangle)$, for $x \in \{0,1\}$ [bases for encoding]
$i := 1$
while $i \leq n$ **do**
Alice chooses $B_i, X_i \in_R \{0,1\}$
Bob chooses $B_i' \in_R \{0,1\}$
Alice prepares a qubit Q_i in state ϕ_{X_i,B_i} and gives it as input to the quantum channel
Bob measures the output Q_i' of the quantum channel w.r.t. basis $\{\phi_{0,B_i'}, \phi_{1,B_i'}\}$ to get Y_i
Alice and Bob communicate B_i and B_i' over the classical channel
if $B_i = B_i'$ **then** $i := i+1$
endwhile
return $(\mathbf{X} = (X_1, \ldots, X_n), \mathbf{Y} = (Y_1, \ldots, Y_n))$

FIG. 12. Prepare-and-measure raw key distribution.

protocol RawKeyDistribution() [entanglement-based]
parameters n [number of signals]; $\phi_{x,0} := |x\rangle$, $\phi_{x,1} := \frac{1}{\sqrt{2}}(|0\rangle + (-1)^x|1\rangle)$, for $x \in \{0,1\}$ [bases for encoding]
$i := 1$
while $i \leq n$ **do**
Alice chooses $B_i \in_R \{0,1\}$
Bob chooses $B_i' \in_R \{0,1\}$
Alice prepares qubits $(\bar{Q}_i, Q_i)$ in state $\frac{1}{\sqrt{2}}(|0\rangle|0\rangle + |1\rangle|1\rangle)$ and gives Q_i as input to the quantum channel
Alice measures $\bar{Q}_i$ w.r.t. basis $\{\phi_{0,B_i}, \phi_{1,B_i}\}$ to get X_i
Bob measures the quantum channel output Q_i' w.r.t. basis $\{\phi_{0,B_i'}, \phi_{1,B_i'}\}$ to get Y_i
Alice and Bob communicate B_i and B_i' over the classical channel
if $B_i = B_i'$ **then** $i := i+1$
endwhile
return $(\mathbf{X} = (X_1, \ldots, X_n), \mathbf{Y} = (Y_1, \ldots, Y_n))$

FIG. 13. Entanglement-based raw key distribution.

This disturbance becomes larger the more information Eve has gained, a fact that is known as the *information-disturbance trade-off* (Fuchs, 1998). In the case of entanglement-based protocols, the bounds on Eve's information can be regarded as an instance of the monogamy of entanglement. It asserts that the stronger that Alice's entanglement with Bob is, the weaker her correlation with Eve is (Coffman, Kundu, and Wootters, 2000; Koashi and Winter, 2004; Terhal, 2004).

A. Protocol replacement

Cryptographic protocols that are optimized for practical use are often not easy to analyze directly. Conversely, protocols that are designed in a way that simplifies their security proofs are usually not easily implementable in practice. For example, building an entanglement-based QKD protocol in practice is technologically more challenging than building a prepare-and-measure scheme. Conversely, the structure of entanglement-based schemes fits more naturally with the known techniques for proving security.

A first step in a security proof for a practical protocol $\pi_{\text{practical}}$ is thus usually to conceive of another protocol $\pi_{\text{theoretical}}$ that is adapted to the proof techniques at hand. One then argues that, for the purpose of the security proof, $\pi_{\text{practical}}$ can be replaced by $\pi_{\text{theoretical}}$, i.e., that the security of $\pi_{\text{practical}}$ is implied by the security of $\pi_{\text{theoretical}}$. A generic way to achieve this is to show that for any possible attack against $\pi_{\text{practical}}$ there is a corresponding attack against $\pi_{\text{theoretical}}$.

For a concrete example, suppose that $\pi_{\text{practical}}$ is the BB84 protocol (Bennett and Brassard, 1984). The protocol follows the generic structure shown in Fig. 11, with a particular raw key distribution procedure similar to that shown in Fig. 12. The protocol prescribes that Alice and Bob proceed in rounds. In each round i, Alice inputs one qubit Q_i to the quantum channel. The qubit encodes a random signal bit X_i with respect to a randomly chosen basis B_i. Bob measures the output Q_i' of the quantum channel with respect to a randomly chosen basis B_i' to obtain a bit Y_i. This is a prepare-and-measure scheme and is in this sense "practical."

The corresponding "theoretical" protocol $\pi_{\text{theoretical}}$ could be an entanglement-based protocol similar to the E91 protocol (Ekert, 1991). This protocol is identical to the previously described BB84 protocol, except that the raw key distribution step is replaced by the procedure shown in Fig. 13. In each round i, Alice creates an entangled state between two qubits $\bar{Q}_i$ and Q_i and sends the latter to Bob, who receives it[31] as Q_i'. Alice and Bob then both select random bases B_i and B_i' and measure their qubits accordingly to obtain bits X_i and Y_i, respectively.

As first shown by Bennett, Brassard, and Mermin (1992), these two protocols $\pi_{\text{practical}}$ and $\pi_{\text{theoretical}}$ are equivalent in terms of their security.[32] Note first that Bob's part of the protocol is the same for $\pi_{\text{practical}}$ and $\pi_{\text{theoretical}}$. To see the correspondence of Alice's part, consider the two bits B_i and X_i together with the qubit Q_i generated by Alice in any round i. It is straightforward to verify that, for both $\pi_{\text{practical}}$ and $\pi_{\text{theoretical}}$, these are described using the same classical-classical-quantum state of the form

$$\rho_{B_i X_i Q_i} = \frac{1}{4} \sum_{b=0}^{1} \sum_{x=0}^{1} |b\rangle\langle b| \otimes |x\rangle\langle x| \otimes |\phi_{x,b}\rangle\langle\phi_{x,b}|. \tag{19}$$

Equation (19) shows, in particular, that, from the viewpoint of an adversary (who may have access to the quantum channel and hence to Q_i), the two protocols are equivalent.

The previously described entanglement-based protocol $\pi_{\text{theoretical}}$ may be further modified to make it even more suitable for security proofs. One such modification concerns the timing of the steps. Instead of running through n rounds, in each of which an entangled qubit pair is created and the qubits measured, one may instead consider a first step in which n entangled qubit pairs $(\bar{Q}_i, Q_i)$ are distributed between Alice and Bob and, rather than being measured directly, are first

[31] Security is also guaranteed if this entangled state is generated by an untrusted third party and distributed to Alice and Bob.

[32] This statement is valid only in the device-dependent setting and does not extend to device-independent security proofs; see Ekert and Renner (2014). For full device-independent security, it is necessary to distribute entanglement.

protocol RawKeyDistribution() [with postponed measurement]
parameters n [number of signals]; $\phi_{x,0} := |x\rangle$, $\phi_{x,1} := \frac{1}{\sqrt{2}}(|0\rangle + (-1)^x|1\rangle)$, for $x \in \{0,1\}$ [bases for encoding]
for $i \in \{1,\ldots,n\}$ **do**
  Alice prepares qubits $(\bar{Q}_i, Q_i)$ in state $\frac{1}{\sqrt{2}}(|0\rangle|0\rangle + |1\rangle|1\rangle)$ and gives Q_i as input to the quantum channel
  Bob stores the quantum channel output Q_i'
endfor
for $i \in \{1,\ldots,n\}$ **do**
  Alice chooses $B_i \in_R \{0,1\}$ and communicates B_i to Bob over the classical channel
  Alice measures $\bar{Q}_i$ w.r.t. basis $\{\phi_{0,B_i'}, \phi_{1,B_i}\}$ to get X_i
  Bob measures Q_i' w.r.t. basis $\{\phi_{0,B_i'}, \phi_{1,B_i}\}$ to get Y_i
endfor
return $(\mathbf{X} = (X_1,\ldots,X_n), \mathbf{Y} = (Y_1,\ldots,Y_n))$

FIG. 14. Entanglement-based raw key distribution with postponed measurement.

stored in quantum memories. Only in a second step do Alice and Bob choose bases $B_i = B_i'$ for each of their qubit pairs and measure them accordingly. This is shown in Fig. 14. An argument similar to the previously given one shows that this change has no impact on the security of the protocol.

B. Raw key distribution and parameter estimation

The first part of the security proof concerns the raw key distribution and the parameter estimation step. For a raw key distribution we consider the subprotocol described in Fig. 14. The parameter estimation is shown in Fig. 15. It essentially calculates an estimate for the fraction η of positions i in which the bit strings $\mathbf{X}$ and $\mathbf{Y}$ differ, i.e., $|X_i - Y_i| = 1$, and returns the value "fail" if this fraction exceeds a given threshold η_0.

To run the raw key distribution and parameter estimation protocol, one needs as initial resources an insecure quantum channel $\mathcal{Q}$ together with an authentic classical channel $\mathcal{A}$, as shown in Fig. 7. The target is a raw key resource $\mathcal{R}$, which can be understood as a weak version of a shared secret key resource, as shown in Fig. 6(b). The resource $\mathcal{R}$ is equipped with a switch controlled by Eve (Portmann, 2017b). If the switch is in position 1, the resource merely outputs $\perp$ to Alice and Bob. If the switch is in position 0, the resource outputs bit strings $\mathbf{X}$ and $\mathbf{Y}$ of length n to Alice and Bob, but at the same time enables Eve to interact with the resource, allowing her to gain information E. Information E is bounded by a secrecy condition, which may be expressed as follows in terms of a

protocol ParameterEstimation$(\mathbf{X},\mathbf{Y})$
parameters s [sample size]; η_0 [threshold]
Alice chooses a subset $S \subset_R \{1,\ldots,n\}$,
  with $n = |\mathbf{X}|$ [length of $\mathbf{X}$] and $s = |S|$ [size of S]
Alice communicates $\{(i, X_i) : i \in S\}$ over the classical channel
Bob computes $\eta = \frac{1}{s}\sum_{i\in S}|X_i - Y_i|$
if $\eta \le \eta_0$ **then return** ok **else return** fail

FIG. 15. Parameter estimation.

lower bound t on the smooth min-entropy (Renner, 2005) of Alice's output $\mathbf{X}$ conditioned on E:

$$H^{\varepsilon}_{\min}(\mathbf{X}|E) \ge t. \tag{20}$$

In Eq. (20) $\varepsilon > 0$ is a small parameter that will contribute additively to the failure probability of the protocol. The choice of this particular measure for entropy will be relevant for the further proof steps to follow, especially privacy amplification. Intuitively, one may think of $H^{\varepsilon}_{\min}(\mathbf{X}|E)$ as the minimum number of bits that can be extracted from $\mathbf{X}$ that are uniform and uncorrelated to E, except with probability ε.

The desired statement is that running the raw key distribution protocol followed by the parameter estimation protocol on $\mathcal{Q}$ and $\mathcal{A}$ constructs the raw key resource $\mathcal{R}$ for appropriately chosen parameters. One may view this as the core of security proofs in QKD. It shows that a criterion on the statistics of the data $\mathbf{X}$ and $\mathbf{Y}$ measured by Alice and Bob, as tested by the parameter estimation protocol, is sufficient to imply a certain level of secrecy of $\mathbf{X}$ toward Eve.

To illustrate the idea behind the argument, we focus for the moment on collective attacks; see Sec. III.C.2. Under this assumption, each of the qubit pairs $(\bar{Q}_i, Q_i')$ held by Alice and Bob when they execute the raw key distribution protocol of Fig. 14 prior to the measurement is in the same state $\rho_{\bar{Q}_i,Q_i'}$. Recall, however, that the second qubit Q_i' is what Bob received. Since Eve may corrupt the quantum communication channel, it is not guaranteed that this qubit coincides with the qubit Q_i that Alice sent. The state $\rho_{\bar{Q}_i,Q_i'}$ may thus be different than the entangled state $(1/\sqrt{2})(|0\rangle|0\rangle + |1\rangle|1\rangle)$ that Alice prepared.

To gain some intuition, it may be useful to consider the special case where the threshold in the subprotocol for parameter estimation is small, say, even $\eta_0 = 0$. If the subprotocol returns the value "ok," this means that the bit strings $\mathbf{X}$ and $\mathbf{Y}$ largely coincide. This yields a constraint on the state $\rho_{\bar{Q}_i,Q_i'}$, namely, that if both Alice and Bob measure it with respect to the basis $\{|0\rangle, |1\rangle\}$ or with respect to the basis $\{(1/\sqrt{2})(|0\rangle \pm |1\rangle)\}$, they obtain identical outcomes, except with some small probability that is due to the finite sample size used for parameter estimation.

It is now straightforward to verify that the only states $\rho_{\bar{Q}_iQ_i'}$ that can pass the test with $\eta_0 = 0$ are those that are close to the pure state $(1/\sqrt{2})(|0\rangle|0\rangle + |1\rangle|1\rangle)$ that Alice prepared. Next one may consider the joint state $\rho_{\bar{Q}_iQ_i'E}$ that includes Eve. But because the state of the first two qubits is almost pure, one can conclude that this state must be of the form

$$\rho_{\bar{Q}_iQ_i'E} \approx \rho_{\bar{Q}_iQ_i'} \otimes \rho_E. \tag{21}$$

That is, Eve's information E is almost uncorrelated to $\bar{Q}_i$ and Q_i'. But because each of the n bits X_i of $\mathbf{X}$ is obtained from a measurement of $\bar{Q}_i$, it is as well almost uncorrelated to E. This proves that each bit X_i is almost uniformly random and independent of E. The smooth min-entropy of the entire sequence $\mathbf{X}$ of bits is thus almost maximal, i.e., $H^{\varepsilon}_{\min}(\mathbf{X}|E) \approx n$.

If instead of $\eta_0 = 0$ one inserts an arbitrary value for the tolerated noise tolerance η_0, which is also known as the quantum bit error rate (QBER), a refinement of the argument that we just sketched gives (Renner, 2005; Renner, Gisin, and Kraus, 2005)

$$H^{\varepsilon}_{\min}(\mathbf{X}|E) \geq n[1 - h(\eta_0)] + O(\sqrt{n}), \tag{22}$$

where $h(x) = -x\log_2(x) - (1-x)\log_2(1-x)$ denotes the binary entropy function.

We also note that the argument can be adapted to the case of device-independent security. In this case the parameter estimation tests whether the outcome statistics of Alice and Bob violates a Bell inequality. The lower bound on the entropy then depends on the degree of this violation; see Acín *et al.* (2007) for the example of the Clauser-Horne-Shimony-Holt (CHSH) Bell inequality (Clauser *et al.*, 1969).

The assumption of collective attacks is necessary to sensibly talk about the state $\rho_{\bar{Q}_i Q'_i}$ of the individual systems. However, there are no good reasons why an adversary should be restricted to such attacks; see Sec. IV.A. Modern security proofs therefore usually consist of an additional step, in which it is shown that general attacks cannot be more powerful than collective attacks.

There are various techniques to achieve this. The most widely one used to date is based on the exponential de Finetti theorem (Renner, 2005, 2007; Renner and Cirac, 2009). The theorem states that if a state over many subsystems, such as $\rho_{\bar{Q}_1 Q'_1 \cdots \bar{Q}_n Q'_n}$, is symmetric under reorderings (i.e., the state remains the same if one permutes the subsystems $\bar{Q}_i Q'_i$), then it is well approximated by a mixture of i.i.d. states, i.e., states of the form $\rho_{\bar{Q}_1 Q'_1} \otimes \cdots \otimes \rho_{\bar{Q}_n Q'_n}$. The latter corresponds to the structure that one has if one assumes collective attacks.

To apply the exponential de Finetti theorem, it is sufficient to argue that the rounds of the protocol, in which the individual signals are sent, could be reordered arbitrarily. As in the example of the previously described BB84 protocol, this is the case for most protocols that have been proposed in the literature. Notable exceptions are the coherent one-way protocol (Stucki *et al.*, 2005) and the differential phase shift protocol (Inoue, Waks, and Yamamoto, 2002), where information is encoded in the correlations between signals.

Another method, which is related to the de Finetti theorem, is the postselection technique (Christandl, König, and Renner, 2009). Like the former, it can be used to lift security proofs against collective attacks to security proofs against general attacks, provided that the protocol satisfies the previously described symmetry assumptions.

Under certain conditions, it is also possible to establish bounds of the form of Eq. (22) directly for general attacks, i.e., without first restricting to collective attacks. This is the case for the approaches presented by Christandl, Renner, and Ekert (2004) and Renner, Gisin, and Kraus (2005), which are both applicable to the device-dependent setting, as well as the techniques proposed by Tomamichel and Renner (2011) and Tomamichel *et al.* (2012), which include semi–device independent scenarios, and by Reichardt, Unger, and Vazirani (2013), Miller and Shi (2014), and Vazirani and Vidick (2014), which apply to particular device-independent protocols.

The most recent approach to directly prove security against general attacks relies on the entropy accumulation theorem (Dupuis, Fawzi, and Renner, 2020). This approach, in contrast to methods based on the de Finetti theorem, gives rather tight min-entropy bounds even when the number n of protocol rounds is relatively small. It is furthermore applicable to the semi–device independent and the device-independent setting (Arnon-Friedman, Renner, and Vidick, 2019), which are discussed in Sec. VI.

C. Information reconciliation

The goal of information reconciliation is to ensure that Alice and Bob have the same (raw) key. The most common way to achieve this is to regard Alice's bit string $\mathbf{X}$ as the key, and to let Bob infer this key from the information $\mathbf{Y}$ that he has. To this end, Alice sends partial information about $\mathbf{X}$ to Bob over the classical channel.

The protocol shown in Fig. 16 uses as resources a raw key $\mathcal{R}$, as described in Sec. V.B, as well as, again, an authentic classical communication channel $\mathcal{A}$. Its purpose is to generate a weak key resource $\mathcal{R}'$, which provides a guarantee of the form of Eq. (20) on the secrecy of the key, and, in addition, ensures that Alice and Bob's values $\mathbf{X}$ and $\mathbf{X}'$ are identical.

We note that information reconciliation is a purely classical subprotocol. It is also largely independent of the other parts of the QKD protocol, and hence works in both the device-dependent and the device-independent setting. The choice of the coding scheme, i.e., the functions enc and dec that the protocol invokes, merely depends on the noise model. The latter describes how Alice's and Bob's inputs to the protocol $\mathbf{X}$ and $\mathbf{Y}$ are correlated with each other.

The noise model is most generally specified in terms of a joint probability distribution of $\mathbf{X}$ and $\mathbf{Y}$. The coding scheme must then be chosen such that

$$\Pr\{\mathrm{dec}[\mathrm{enc}(\mathbf{X}, \mathbf{Y})] = \mathbf{X}\} \geq 1 - \varepsilon. \tag{23}$$

The parameter $\varepsilon > 0$ bounds the failure probability of the subprotocol and will hence, like the parameter ε used in the previous step, contribute additively to the total failure probability of the QKD protocol. Furthermore, to maintain as much secrecy as possible for $\mathbf{X}$, the function enc should be chosen such that $C = \mathrm{enc}(\mathbf{X})$ does not reveal too much information about $\mathbf{X}$. (Recall that the classical channel is accessible to Eve, so she may get ahold of C.) This may be achieved by making C as small as possible. It can be shown using classical techniques

1. **protocol** InformationReconciliation($\mathbf{X}, \mathbf{Y}$)
2. **parameters** enc, dec [coding scheme]
3. Alice sends $C = \mathrm{enc}(\mathbf{X})$ over the classical channel
4. Bob computes $\mathbf{X}' = \mathrm{dec}(C, \mathbf{Y})$
5. **return** $(\mathbf{X}, \mathbf{X}')$

FIG. 16. Information reconciliation.

from information theory that any coding scheme that satisfies Eq. (23) requires a communication C of

$$k \geq H^{\varepsilon}_{\max}(\mathbf{X}|\mathbf{Y}) \text{ bits}, \tag{24}$$

where $H^{\varepsilon}_{\max}$ denotes the smooth max-entropy (Renner and Wolf, 2005). Furthermore, there are coding schemes that saturate this bound (up to a small additive constant).

In the case of an i.i.d. noise model, $H^{\varepsilon}_{\max}(\mathbf{X}|\mathbf{Y})$ is approximated by the Shannon entropy up to terms of the order of $\sqrt{n}$, where n is the length of $\mathbf{X}$. For a protocol such as BB84, which uses single qubits, and assuming that the QBER is η_0, one thus has

$$k \approx nh(\eta_0) + O(\sqrt{n}). \tag{25}$$

Letting E be the initial information that Eve has about the raw key $\mathbf{X}$ before information reconciliation, the secrecy after information reconciliation with communication C consisting of k bits is given by

$$H^{\varepsilon}_{\min}(\mathbf{X}|EC) \gtrapprox H^{\varepsilon}_{\min}(\mathbf{X}|E) - k - O(1). \tag{26}$$

Hence, for an optimal information reconciliation protocol we have

$$H^{\varepsilon}_{\min}(\mathbf{X}|EC) \gtrapprox H^{\varepsilon}_{\min}(\mathbf{X}|E) - H_{\max}(\mathbf{X}|\mathbf{X}') - O(\sqrt{n}). \tag{27}$$

In particular, for the case of the BB84 protocol, we get

$$H^{\varepsilon}_{\min}(\mathbf{X}|EC) \geq n[1 - 2h(\eta_0)] - O(\sqrt{n}). \tag{28}$$

As is clear from Eq. (26), the amount of secrecy that is left after information reconciliation depends on the amount k of communication required. The design of coding schemes (enc, dec) that optimize this parameter is a main subject of classical information theory (Cover and Thomas, 2012). While the bound in Eq. (24) can already be saturated with randomly constructed encoding functions, a main challenge is to develop schemes for which the encoding and decoding functions are efficiently computable (Leverrier *et al.*, 2008; Elkouss *et al.*, 2009; Elkouss, Martinez-Mateo, and Martin, 2011; Jouguet and Kunz-Jacques, 2014).

While the information reconciliation protocol of Fig. 16 invokes only one-way communication from Alice to Bob, one may also consider two-way schemes. In fact, the first proposals for QKD implementations used a procedure to correct errors that required multiple rounds of communication between Alice and Bob (Bennett, Bessette *et al.*, 1992).[33] Furthermore, one may also include advantage distillation (Maurer, 1993). Here the idea is that Alice and Bob group their data into small blocks. They then try to distinguish blocks that are likely to contain few errors from those that are likely to contain many errors. The ones with many errors are then discarded. It has been shown that this technique can be advantageous relative to standard error correction (Gottesman and Lo, 2003; Renner, 2005; Tan, Lim, and Renner, 2020).

1. **protocol** PrivacyAmplification$(\mathbf{X}, \mathbf{X}')$
2. **parameters** $\{\mathrm{ext}_s\}_{s\in\mathcal{S}}$ [randomness extractor]
3. Alice chooses $S \in_R \mathcal{S}$ and sends it over the classical channel
4. Alice computes $K = \mathrm{ext}_S(\mathbf{X})$
5. Bob computes $K' = \mathrm{ext}_S(\mathbf{X}')$
6. **return** (K, K')

FIG. 17. Privacy amplification.

D. Privacy amplification

The aim of privacy amplification is to turn the weakly secret key $\mathbf{X}$, which after information reconciliation is known to Alice and Bob, into a strong secret key K, i.e., a bit string that is essentially uniform and independent of the information held by an adversary (Bennett, Brassard, and Robert, 1988; Bennett *et al.*, 1995). This is typically achieved with a protocol as in Fig. 17. Apart from the weak key resource $\mathcal{R}$, which satisfies a secrecy bound of the form of Eq. (20) and which is assumed to output the same string $\mathbf{X}$ to Alice and Bob, the protocol requires an authentic communication channel $\mathcal{A}$. From these resources, the protocol constructs a secret key resource as shown in Fig. 6(b).

The protocol makes use of a randomness extractor (Zuckerman, 1990; Shaltiel, 2004). This is a family of functions ext_s parametrized by a seed $s \in \mathcal{S}$, which take as input a bit string such as $\mathbf{X}$ and output a bit string of a fixed length ℓ. In the classical literature, a strong (k, ε) extractor is defined by the property that, for any input $\mathbf{X}$ whose min-entropy satisfies the lower bound $H_{\min}(\mathbf{X}) \geq k$, the output $\mathrm{ext}_s(\mathbf{X})$ is ε close to uniform. More precisely, the expectation over a randomly chosen seed $s \in \mathcal{S}$ of the variational distance between the distribution of the output $\mathrm{ext}(\mathbf{X})$ and a uniform string U of ℓ bits must be upper bounded by ε,

$$\mathrm{Exp}_s[D(P_{\mathrm{ext}_s(\mathbf{X})}, P_U)] \leq \varepsilon. \tag{29}$$

Equation (29), however, does not take into account the quantum nature of information that an adversary may have about $\mathbf{X}$ (König, Maurer, and Renner, 2005; Gavinsky *et al.*, 2007). It is hence not sufficient for use in the context of quantum key distribution, unless one restricts to security against individual attacks, which corresponds to forcing the adversary to store classical information only; see Sec. IV.A.

To prove general security, it is necessary to demand that the randomness extractor $\{\mathrm{ext}_s\}_{s\in\mathcal{S}}$ be quantum proof for parameters k and ε, as previously shown. This means that, for any $\mathbf{X}$ and any quantum system E such that $H_{\min}(\mathbf{X}|E) \geq k$, one has

$$\mathrm{Exp}_s[D(\rho_{\mathrm{ext}_s(\mathbf{X})E}, \rho_U \otimes \rho_E)] \leq \varepsilon. \tag{30}$$

[33] Despite its two-way nature, the particular method proposed by Bennett, Bessette *et al.* (1992) did not achieve the previously described information-theoretic boundse. It was realized only later by Bennett, Brassard *et al.* (1992), in the context of oblivious transfer, that one-way error correction is sufficient and can be made asymptotically optimal.

Note that this criterion refers to min-entropy $H_{\min}(\mathbf{X}|E) = H^{\varepsilon'}_{\min}(\mathbf{X}|E)$ with smoothness parameter $\varepsilon' = 0$. However, a straightforward application of the triangle inequality for the distance between states implies that a corresponding criterion also holds if $\varepsilon' > 0$ (Renner, 2005).

A number of constructions for quantum-proof extractors have been proposed in the literature (Renner, 2005; Fehr and Schaffner, 2008; König and Terhal, 2008; Ben-Aroya and Ta-Shma, 2012; De *et al.*, 2012; Mauerer, Portmann, and Scholz, 2012; Berta, Fawzi, and Scholz, 2017). In the context of QKD, the most widely used extractors are based on two-universal hashing (Carter and Wegman, 1979; Wegman and Carter, 1981). As shown by Renner (2005), Renner and König (2005), and Tomamichel *et al.* (2010), these can achieve an output length of $\ell = k - 2\log_2(1/\varepsilon)$ while still being quantum-proof (k, ε) extractors. Using them within the protocol of Fig. 17 generates a key of length

$$\ell = H^{\varepsilon}_{\min}(\mathbf{X}|EC) - O(1), \tag{31}$$

with a failure probability of the order of ε. Combining this with the previously discussed results, with optimal information reconciliation and privacy amplification it is possible to generate a key of length

$$\ell = H^{\varepsilon}_{\min}(\mathbf{X}|E) - H_{\max}(\mathbf{X}|\mathbf{Y}) - O(1). \tag{32}$$

In particular, in the case of the BB84 protocol, we obtain

$$\ell = n[1 - 2h(\eta_0)] - O(\sqrt{n}), \tag{33}$$

where η_0 is the QBER. The asymptotic key rate is thus $1 - 2h(\eta_0)$.

E. Other approaches to proving security

The previously described generic security proof follows the approach proposed by Renner (2005). It is sometimes termed information theoretic, as its core part consists of bounds on entropic quantities, such as Eq. (22). Such bounds were first proposed by Devetak and Winter (2005). They were further developed by Renes and Renner (2012) and used by Kraus, Gisin, and Renner (2005) and Renner, Gisin, and Kraus (2005); see also Christandl *et al.* (2007) for related work. However, as already mentioned, there are a variety of other proof strategies.

Early proofs (Lo and Chau, 1999; Shor and Preskill, 2000) used a reduction to the problem of entanglement distillation. For this, one rearranges the key distribution protocol such that all measurements are postponed to the last step. If one now omits these final measurements, Alice and Bob end up with correlated quantum registers rather than classical keys. One may thus regard the protocol as an entanglement distillation protocol (Bennett, Bernstein *et al.*, 1996; Bennett, Brassard *et al.*, 1996) and prove that the registers held by Alice and Bob are almost maximally entangled. If this is the case, then, by the monogamy of entanglement, the information in these registers is uncorrelated to Eve, and hence secret.[34]

This approach may be more generally understood as follows. Assuming that Alice's and Bob's start with quantum correlation stored in individual qubits equipped with a computational basis, the entanglement distillation protocol can be regarded as a quantum error correction scheme (Calderbank and Shor, 1996; Steane, 1996) that corrects for both bit and phase flip errors. The correction of bit flip errors ensures that Alice and Bob end up with the same key. The correction of phase flip errors ensures that the two registers are not only classically correlated but also maximally entangled. Since, as indicated earlier, the latter implies secrecy, one can understand the correction of phase flip errors as a kind of privacy amplification (Renes, 2013).

The technique was originally used to prove the security of the BB84 protocol, including variants with imperfect devices (Gottesman *et al.*, 2004), but can also be applied to other quantum key distribution protocols (Tamaki, Koashi, and Imoto, 2003; Koashi, 2004; Boileau *et al.*, 2005). While the correspondence to entanglement distillation requires that error correction and privacy amplification be treated as a single quantum error correction step, it is under certain conditions possible to achieve a separation in a way similar to the previously described modular description (Lo, 2003). Furthermore, as shown by Horodecki *et al.* (2008), the method also works if the registers of Alice and Bob merely contain bound entanglement, i.e., entanglement from which no maximally entangled states can be distilled (Horodecki, Horodecki, and Horodecki, 1998).

A somewhat related strategy, proposed originally by Mayers (2001), is the use of complementarity (Koashi, 2009). Specifically, one uses the fact that if Alice is able to accurately predict the outcomes of a measurement in one basis, say, the computational basis, then by the uncertainty principle any predictions for the outcomes of measurements in a complementary basis, e.g., the diagonal basis in the case of single qubits, will be inaccurate. This technique has been refined in a series of works and made applicable to the study of finite-size effects (Hayashi and Tsurumaru, 2012; Tomamichel *et al.*, 2012) and to measurement device–independent cryptography (Tamaki *et al.*, 2012). The complementarity approach is also related to the use of entropic uncertainty relations (Berta *et al.*, 2010; Tomamichel and Renner, 2011; Coles *et al.*, 2017).

VI. ALTERNATIVE MODELING OF QKD

Thus far we have discussed QKD as protocols that start with an insecure quantum channel and an authentic classical channel and generate, as the desired ideal resource, a key of fixed length. In this section we discuss other variants of QKD protocols, where these resources are chosen differently. In Sec. VI.A we consider an ideal key resource with adaptive

[34]The following is a quantitative version of this statement. If the entanglement distilled by Alice and Bob has fidelity F to a maximally entangled state, then it follows from Theorem 1 of Fuchs and Graaf (1999) that the corresponding security parameter ε according to Eq. (8) is bounded by $\epsilon \leq \sqrt{1 - F^2}$.

key length. In Sec. VI.B we discuss protocols that use a source of entanglement instead of an insecure quantum channel. In Sec. VI.C we show how to model a situation in which no perfect randomness is available. In Sec. VI.D we model device-independent QKD. Relaxations of this known as semi–device independence are discussed in Sec. VI.E. Finally, in Sec. VI.F we consider adversaries that have no quantum memory.

A. Adaptive key length

For a protocol to construct the shared secret key resource of Fig. 6(b), it must either abort or produce a key of a fixed length. A more practical protocol could adapt the secret key length to the noise level of the quantum channel. This provides the adversary with the functionality to control the key length (not only whether or not it gets generated) and can be modeled by allowing the key length to be input at Eve's interface of the ideal key resource, as illustrated in Fig. 18.

Such an ideal resource was considered by Ben-Or *et al.* (2005) and Hayashi and Tsurumaru (2012). The reduction from the corresponding security definition in AC to a trace distance criterion still goes through. But instead of Eq. (7) we get

$$\sum_m p_m D(\rho^m_{KE}, \tau^m_K \otimes \rho^m_E) \leq \varepsilon, \tag{34}$$

where p_m is the probability of obtaining a key of length m, ρ^m_{KE} is the joint state of the key and Eve's system conditioned on the key having length m, and τ^m_K is a fully mixed state of dimension 2^m.

B. Source of entanglement

In contrast to prepare-and-measure protocols, entanglement-based protocols (Ekert, 1991; Bennett, Brassard, and Mermin, 1992) use a source of entanglement instead of a quantum communication channel. It is also fairly standard in security proofs to first transform a given prepare-and-measure protocol into an entanglement-based one, then prove the security of the latter (Shor and Preskill, 2000). In Fig. 19 we draw a system consisting of a QKD protocol $(\pi^{\mathrm{QKD}}_A, \pi^{\mathrm{QKD}}_B)$, the authentic channel $\mathcal{A}$, and a source $\mathcal{E}$ of entangled states that may be controlled by Eve. To specify the completeness property, we also consider a source of entanglement $\mathcal{E}'$ that produces a fixed bipartite entangled state instead of allowing Eve to decide.

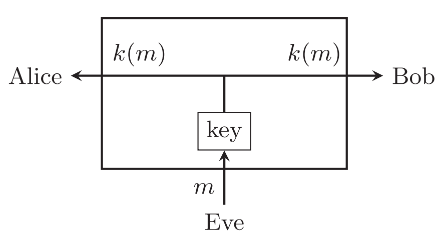

FIG. 18. A secret key resource with adaptive key length. This resources allows Eve to choose the length m of the final key k, which is then output at Alice's and Bob's interfaces.

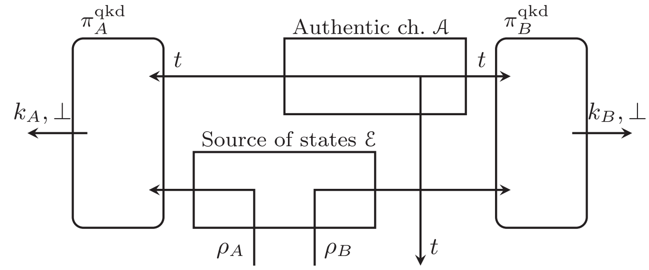

FIG. 19. A real QKD system that uses a source of entangled states. Instead of having access to an insecure channel as in Fig. 7(a), Alice and Bob use a source of entanglement $\mathcal{E}$ that is controlled by Eve. This means that Eve may generate an arbitrary state ρ_{ABE}, of which the A register goes to Alice and the B register goes to Bob.

The reduction from the AC security definition to the trace distance criterion described in Sec. III.B works here too, with the source of entanglement replacing the insecure channel, thus resulting in the same conditions for ε secrecy and ε correctness.

One can also show that any protocol designed for a distributed source of entanglement can be transformed into one where a state is prepared locally and sent over an insecure channel. To explain this, we first decompose Alice's QKD protocol into two parts. In the first she carries out a subprotocol α that performs a measurement $\mathbb{M}^a = \{M^a_x\}_x$ on the state received from the source of entangled states, where $\mathbb{M}^a$ is chosen with some probability p_a from a set $\{\mathbb{M}^a\}_a$. The second part consists of the rest of her QKD protocol. We illustrate this in Fig. 20.

We now need to argue that there is a converter γ that constructs $\alpha\mathcal{E}$ from an insecure channel $\mathcal{Q}$ and $\alpha\mathcal{E}'$ from a noiseless channel $\mathcal{Q}'$. For this, we must establish the two following conditions.

(i) There is a simulator σ_E such that

$$\gamma\mathcal{Q} = \alpha\mathcal{E}\sigma_E.$$

(ii) The following equality holds:

$$\gamma\mathcal{Q}' = \alpha\mathcal{E}'.$$

Once we have established these conditions, it follows immediately from the composition theorem of the AC framework

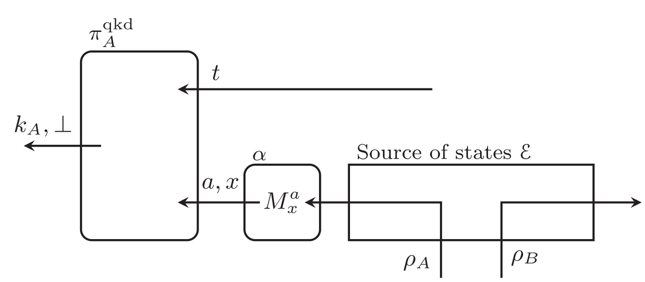

FIG. 20. We split Alice's part of an entanglement-based QKD protocol into two parts, the measurement of the incoming states (denoted by α) and the rest of the protocol (denoted by π^{QKD}_A).

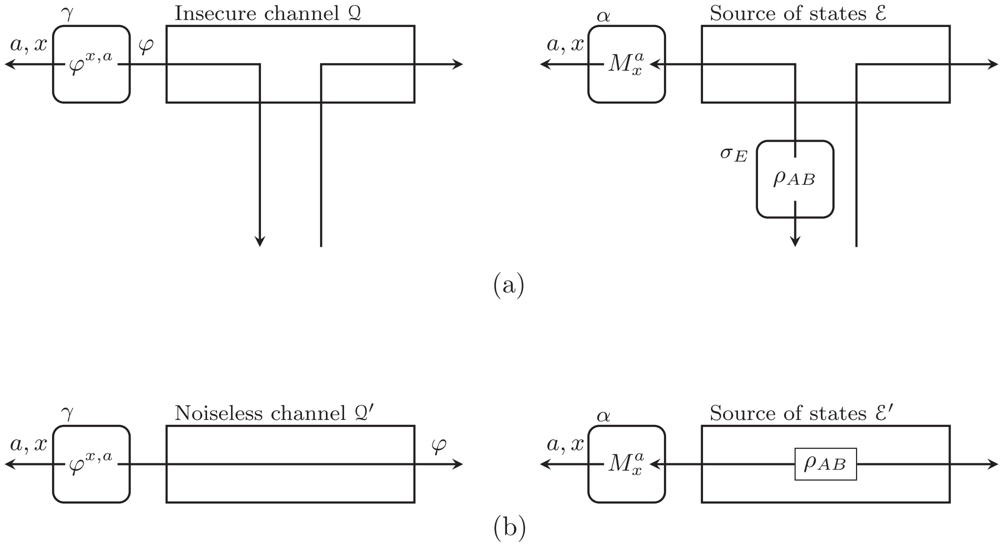

FIG. 21. Pictorial proof for the security of the construction of $\alpha\mathcal{E}$ from $\mathcal{Q}$ and $\alpha\mathcal{E}'$ from $\mathcal{Q}'$. Any protocol that is designed to run with a source of entangled states $\mathcal{E}$ and that measures the incoming states on Alice's side as does α can be equivalently used with an insecure channel $\mathcal{Q}$ and a converter γ that generates the states to be sent on the channel. (a) When modeling soundness, the adversary can modify the messages on the insecure channel $\mathcal{Q}$. The simulator σ_E generates the entangled state ρ_{AB} that is expected from a nonadversarial source of entangled states and outputs the B part at the outer interface, making the systems on the left and right indistinguishable. (b) When modeling completeness, the source of entanglement $\mathcal{E}'$ prepares the state ρ_{AB}. The systems on the left and right are indistinguishable.

(Maurer and Renner, 2011) that any QKD protocol that is sound when using $\alpha\mathcal{E}$ and complete when using $\alpha\mathcal{E}'$ is also sound and complete when using $\gamma\mathcal{Q}$ and $\gamma\mathcal{Q}'$, respectively.

Let ρ_{AB} be the bipartite entangled state that is generated by $\mathcal{E}'$. Let $\tilde{\varphi}_B^{x,a} \coloneqq \mathrm{tr}_A[M_x^a \rho_{AB} (M_x^a)^\dagger]$, $p_{x|a} \coloneqq \mathrm{tr}\tilde{\varphi}_B^{x,a}$, and $\varphi_B^{x,a} \coloneqq \tilde{\varphi}_B^{x,a}/p_{x|a}$. We define the converter γ to prepare the state $\varphi_B^{x,a}$ with probability $p_a p_{x|a}$, which it sends on the insecure channel. Furthermore, we define the simulator σ_E to prepare ρ_{AB}, input the A part on the entanglement resource for Alice, and output the B part at the outer interface. It is then straightforward to check from Fig. 21 that this satisfies the previously described conditions (i) and (ii).

C. Imperfect randomness

QKD protocols usually assume that the honest parties have arbitrary access to perfect random numbers. This,

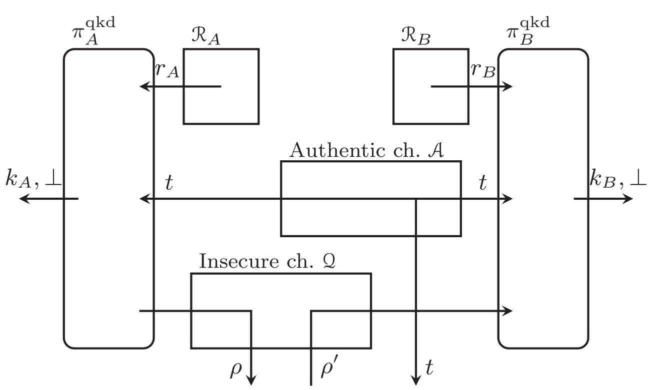

FIG. 22. A real QKD system with a deterministic protocol $(\pi_A^{\mathrm{QKD}}, \pi_B^{\mathrm{QKD}})$ and explicit sources of randomness $\mathcal{R}_A$ and $\mathcal{R}_B$.

however, is never the case in practice. A more realistic model of a QKD system would consider randomness as a resource that is available in limited and imperfect quantities to Alice and Bob. The real QKD setting drawn in Fig. 7 needs to be changed to take this into account. In Fig. 22 we depict a QKD protocol that, in addition to the insecure quantum channel and authentic classical channel, has access to resources producing local randomness $\mathcal{R}_A$ and $\mathcal{R}_B$ at Alice's and Bob's interfaces, respectively. A different model of randomness resources might also provide some partial quantum information about the randomness to the eavesdropper. For simplicity, however, we chose to draw the simpler case in which $\mathcal{R}_A$ and $\mathcal{R}_B$ have an empty interface for the dishonest party.

In such a setting, the converters π_A^{QKD} and π_B^{QKD} are deterministic systems. A QKD protocol would then construct an ideal key resource given access to these three resources. It remains an open problem to minimize the assumptions on the sources of randomness in QKD. Recent results on device-independent randomness amplification (Colbeck and Renner, 2012) show that under certain minimal assumptions[35] about the workings of an unknown quantum system, one can transform a single public weak source of randomness into a fully private random source (Chung, Shi, and Wu, 2014; Brandão *et al.*, 2016; Kessler and Arnon-Friedman, 2020). Alternatively, if two or more sources of weak randomness are available to a player (under certain strict conditions on the correlations between these different sources), these can be combined to obtain approximately uniform randomness

[35]One generally has to assume that no messages leave or enter the quantum devices unless authorized by the protocol. Some papers make additional assumptions to simplify the protocols and proofs.

(Chung, Li, and Wu, 2014; Arnon-Friedman, Portmann, and Scholz, 2016). Composing this with a standard QKD protocol would allow secret keys to be distributed when only weak randomness is available to the honest parties.

D. Device-independent QKD

In this review we have thus far always considered scenarios for which it is assumed that the players have trusted quantum devices that work exactly according to their specifications. For instance, if a player instructs the device to generate a $|0\rangle$ state, then it is assumed that the device generates precisely this state. This assumption, however, is not met in any actual implementation with realistic devices, as these are never perfect. Indeed, there have been numerous demonstrations of successful attacks against implementations of quantum cryptographic protocols that exploit deviations of the devices' functionality from the specifications, as discussed in Sec. IV. This problem cannot be solved only by a more careful design of the devices, for it appears to be impossible to guarantee their perfect working under all possible environmental conditions.

A theoretical solution to this problem is to devise protocols whose security does not rely on the assumption that devices are perfect. Ideally, they should provide security guarantees even if the devices are untrusted, meaning that their behavior may deviate arbitrarily from the specification. Using quantum devices, this is possible (with certain caveats to be described). The idea is to use a phenomenon called Bell nonlocality (Bell, 1964); see Scarani (2013) and Brunner *et al.* (2014) for articles on the topic. The subfield of cryptography that studies the use of nonlocality to design protocols that work with untrusted devices is termed device-independent cryptography.

In a nutshell, a Bell inequality is a bound on the probability of observing certain values in an experiment involving measurements of two isolated (and hence noncommunicating) systems. The bound characterizes classical locality: it cannot be violated if the two isolated systems are described by classical physics. However, the bound can be violated by measurements on entangled quantum systems. One of the most commonly used Bell inequalities is the CHSH inequality (Clauser *et al.*, 1969). It states that if two players each hold noncommunicating systems and each performs one out of two binary measurements chosen uniformly at random on their respective system, where the choice of the measurement is given by $x, y \in \{0,1\}$ and the outcome is given by $a, b \in \{0,1\}$, respectively, then the probability[36] that $xy = a \oplus b$ should be less than or equal to $3/4$. But if the systems are quantum, it is possible to observe this outcome with probability up to ≈ 0.85: this is achieved if the systems are in a perfectly entangled state and the players perform an optimal measurement.

[36]An alternative formulation of the inequality is $|E(0,0) + E(0,1) + E(1,0) - E(1,1)| \leq 2$, where $E(x,y)$ is the expected value of the product of the outcomes of the systems when measured with settings x and y, respectively, and the outcomes are values in $\{-1,+1\}$.

An observation of a violation of a Bell inequality implies that the measurement outcomes contain some genuine randomness (Colbeck, 2006; Pironio *et al.*, 2010; Acín, Massar, and Pironio, 2012; Colbeck and Renner, 2012), even conditioned on the knowledge of the person who set up and programmed the devices used in the experiment. The only assumptions are that no information other than the measurement result leaves the devices and that these devices never fall into the hands of an adversary, since their internal memory may contain a copy of the measurement outcomes. This randomness may then be used to generate uniform random numbers (Vazirani and Vidick, 2012; Chung, Shi, and Wu, 2014; Miller and Shi, 2014; Brandão *et al.*, 2016; Kessler and Arnon-Friedman, 2020) or a shared secret key (Barrett, Hardy, and Kent, 2005; Pironio *et al.*, 2009; Vazirani and Vidick, 2014; Arnon-Friedman *et al.*, 2018; Arnon-Friedman, Renner, and Vidick, 2019).

For a review of different results and techniques in device-independent cryptography, see Ekert and Renner (2014). In this section we show how to model device-independent quantum key distribution (DI QKD) in the AC framework. It then follows from the composition theorem of AC that the resulting key can be safely used in applications requiring one.

The converters π_A^{QKD} and π_B^{QKD} modeling Alice's and Bob's parts of the protocol in Sec. III are systems that generate quantum states and perform measurements. In DI QKD, exactly these operations cannot be trusted. Instead, the DI protocol $(\pi_A^{\mathrm{DI\ QKD}}, \pi_B^{\mathrm{DI\ QKD}})$ will involve only classical operations. Everything quantum is moved into a resource, a device $\mathcal{D}$. The honest players can send bits to these devices and receive bits back from them. This corresponds to choosing a measurement $x, y \in \{0,1\}$ and receiving the outcome $a, b \in \{0,1\}$, as just described. The adversary is permitted to "program" these devices by providing some initial state ρ as input. Depending on the model, Eve may be allowed to provide further inputs to the device at some later point, such as to provide more Einstein-Podolsky-Rosen (EPR) pairs so that the device may continue running. The corresponding real world is drawn in Fig. 23. The ideal world will be identical to that of standard QKD since we wish to construct the same key resource, i.e., Fig. 8.

When one applies Definition 1, this means that the protocol $(\pi_A^{\mathrm{DI\ QKD}}, \pi_B^{\mathrm{DI\ QKD}})$ constructs $\mathcal{K}$ from $\mathcal{A}$, $\mathcal{D}_A$, and $\mathcal{D}_B$ within ε if

$$\exists\, \sigma_E, \qquad \pi_A^{\mathrm{DI\ QKD}} \pi_B^{\mathrm{DI\ QKD}} (\mathcal{D}_A \| \mathcal{D}_B \| \mathcal{A}) \approx_\varepsilon \mathcal{K}\sigma_E. \tag{35}$$

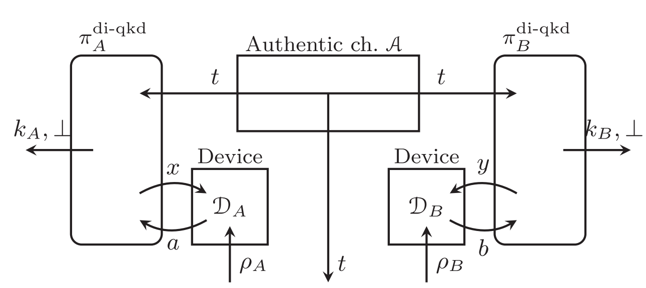

FIG. 23. The real-world setting for a DI-QKD protocol. Eve can program the devices $\mathcal{D}$ but cannot receive any output from them.

Note that we have not specified the behaviors of the devices $\mathcal{D}_A$ and $\mathcal{D}_B$ at all. In fact, we need Eq. (35) to hold for all devices[37] $\mathcal{D}_A$ and $\mathcal{D}_B$. This is exactly the device-independent guarantee, namely, that security holds regardless of how the quantum devices work. Alternatively, one can consider fixed devices $\mathcal{D}_A$ and $\mathcal{D}_B$ that are universal computers and have their program be part of the inputs at the E interface.

As usual, completeness is captured by specific devices $\mathcal{D}'_A$ and $\mathcal{D}'_B$ that work honestly (e.g., they share perfectly entangled states and perform the correct measurements as specified by the protocol), as well as the same honest resources $\mathcal{A}'$ and $\mathcal{K}'$ as in Sec. III. In addition to Eq. (35), we need

$$\pi_A^{\mathrm{DI\ QKD}}\pi_B^{\mathrm{DI\ QKD}}(\mathcal{D}'_A\|\mathcal{D}'_B\|\mathcal{A}') \approx_{\varepsilon'} \mathcal{K}'.$$

The same reduction as for normal QKD goes through, and one can show that Eq. (35) is satisfied if, for all behaviors of the devices (and their inputs), Eqs. (13) and (14) hold.

Note, however, that the construction outlined in this section allows the devices $\mathcal{D}_A$ and $\mathcal{D}_B$ to only be accessed during the protocol. No access is granted after the protocol ends, meaning that we make no security statement about what happens if the devices are reused. It is an open question how to reuse devices in DI cryptography, which we discuss in Sec. IX.A.

Proving the security of device-independent QKD is more challenging than in the device-dependent case. One of the difficulties is that the measurement operators that describe Alice's and Bob's measurements can be arbitrary. In particular, it cannot be assumed that two subsequent measurement outcomes by Bob are obtained by two separate measurement processes. While some of the techniques described in Sec. V, such as entropy accumulation, are still applicable to the device-independent setting, others, like de Finetti–type theorems, are not or must be adapted; see Arnon-Friedman (2018) for details.

E. Semi–device independent QKD

The only assumption made about the devices in DI QKD is that no information leaves these devices unless it is allowed by the protocols; see Sec. VI.D. But achieving the violation of Bell inequalities needed for this is challenging because it requires high detector efficiency and tolerates only low noise on the channel (Brunner *et al.*, 2014). Protocols that are easier to implement can be achieved by making additional assumptions about the quantum devices used by Alice and Bob. These are generally called semi–device independent (semi-DI).

Many different assumptions may be labeled semi-DI. For example, in a one-sided device-independent setting the protocol is DI for Bob but not for Alice (Branciard *et al.*, 2012). One may also assume dimension bounds on the quantum systems generated by untrusted devices such as those discussed by Pawłowski and Brunner (2011). Alternatively, Lim *et al.* (2013) assumed that each use of the devices is causally independent (i.e., the states generated and measurements performed are in product form) to analyze a protocol where the Bell violation is measured locally in Alice's lab, thus avoiding the noise introduced by the channel between Alice and Bob. Similar ideas have been used for protocols other than QKD, such as semi-DI quantum money (Bozzio, Diamanti, and Grosshans, 2019; Horodecki and Stankiewicz, 2020).

One of the most promising forms of semi-DI QKD, which has already been implemented over large distances (Liu *et al.*, 2013; Tang *et al.*, 2014; Pirandola *et al.*, 2015; Yin *et al.*, 2016), is measurement device–independent (MDI) QKD (Braunstein and Pirandola, 2012; Lo, Curty, and Qi, 2012; Ma and Razavi, 2012; Curty *et al.*, 2014). Here one assumes that players can generate the states they desire, but one does not trust measurement devices at all. This model is motivated by the attacks on the detectors, e.g., the time-shift attacks or detector blinding attacks discussed in Sec. IV.C.

To understand how such protocols work, we start with an entanglement-based protocol as in Sec. VI.B, then modify it step by step until we achieve a prepare-and-measure protocol in which all measurements are performed by the adversary. Since the final protocol is as secure as the original one and the original one is secure for all adversaries, the final MDI-QKD protocol is secure for all adversaries as well. In particular, it is secure for adversaries that completely control the measurement apparatus.

In an entanglement-based protocol, Alice and Bob receive the A and B parts of a state ψ_{ABR} and measure these systems in either the computational or diagonal basis, thereby obtaining a raw key. This key is then processed as in a prepare-and-measure protocol; see Secs. III.A.2 and V. If the source gave them a state that is close to a tensor product of EPR pairs, such a protocol is guaranteed to terminate with a shared secret key. Equivalently, the source could generate any of the Bell states and notify Alice and Bob which one it gave them. They then perform bit or phase flips to change it to an EPR pair.

Instead of the source distributing an entangled state, Alice and Bob could each generate an EPR pair AA' or BB', respectively. They then send A' and B' to a third party, Charlie, who measures this in the Bell basis and tells them the measurement outcome. If performed correctly, the AB system will be in a Bell state, and the measurement outcome will tell them which one. By flipping bits or phases, Alice and Bob can turn this into an EPR pair, and continue with the protocol as previously discussed. If Charlie does not perform the correct measurement, then Alice and Bob will end up holding the A and B parts of some unknown state ψ_{ABR}. But this does not compromise security: if it is too far from the expected state, the protocol will simply abort.

Instead of first performing a bit or phase flip and then measuring, Alice and Bob could first measure their systems A and B and then flip the value of the measurement outcome if needed. And instead of generating EPR pairs AA' and BB', then measuring A and B, they could pick the measurement outcome at random, then generate the corresponding reduced state in A' and B', respectively, and send these to Charlie. When they obtain the measurement outcome from Charlie, they flip their bits if needed.

[37] The simulator may depend on these devices, i.e., $\forall \mathcal{D}_A, \mathcal{D}_B, \exists \sigma_E$ such that Eq. (35) holds.

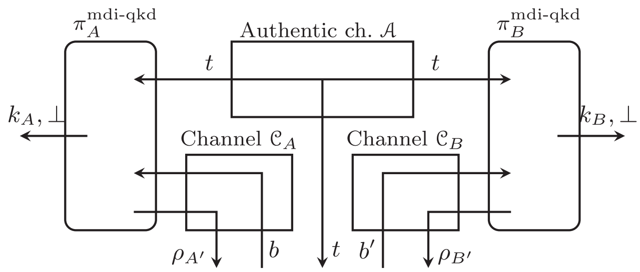

FIG. 24. The real-world setting for a MDI-QKD protocol. The only quantum operations performed by $\pi_A^{\text{MDI QKD}}$ and $\pi_B^{\text{MDI QKD}}$ are the generation of quantum states. The communication resources $\mathcal{C}$ send quantum systems from Alice or Bob to Eve, and classical bits from Eve to Alice and Bob.

The only trusted quantum operations that Alice and Bob need to perform in the protocol just described are generating the states in the systems A' and B'. All measurements have now been delegated to Charlie, who may deviate arbitrarily from the protocol without compromising security.

The real world for such a MDI-QKD protocol is drawn in Fig. 24, where one can see that the converters $\pi_A^{\text{MDI QKD}}$ and $\pi_B^{\text{MDI QKD}}$ do not have any incoming quantum states; i.e., they do not need to perform any measurements.

Security proofs for MDI-QKD protocols can be based on the same techniques as those for fully device-independent protocols, as discussed in Sec. VI.D. The comments on security proofs made in that section thus also apply here.

F. Memoryless adversaries

Thus far we have analyzed different models of QKD in which we have changed the capabilities and resources of the honest players running the protocol. Similar techniques may also be used to model limitations on adversaries. In this section we consider an example of a QKD protocol with an adversary that has no long-term quantum memory and is forced to measure the quantum states exchanged between Alice and Bob during the QKD protocol and to store the classical information.

The insecure channel resource $\mathcal{Q}$ modeled as part of the real QKD system in Fig. 7(a) gives complete control over the states sent on this channel to the adversary. Since this may include storing them and measuring them at a later point, we need to limit the adversary's access to this channel as part of the insecure channel resource. We may thus define a different channel $\tilde{\mathcal{Q}}$ that requires Eve to input some measurement specifications and then obtains the measurement outcome at her interface. The resulting postmeasurement state is then output at Bob's interface.

The previously described model of $\tilde{\mathcal{Q}}$ is merely one possible way one may imagine limiting Eve's access to the states sent during QKD. The result is a change in the requirements of the QKD protocol. Instead of constructing a secure key $\mathcal{K}$ from an authentic channel $\mathcal{A}$ and an insecure channel $\mathcal{Q}$, it is now sufficient if $\mathcal{K}$ can be constructed from $\mathcal{A}$ and $\tilde{\mathcal{Q}}$. Since the accessible information (see Sec. III.C.1) measures the information an adversary has after measuring their quantum states, a QKD protocol with low accessible information satisfies such a construction, namely, $\mathcal{A}\|\tilde{\mathcal{Q}} \to \mathcal{K}$. The accessible information security measure is thus a composable criterion under the assumption that the adversary has such a physical limit on their memory.

Since QKD protocols are secure against general adversaries, as modeled in Sec. III, there does not seem to be much incentive to consider adversaries with limitations on their memory (unlike for certain two-party protocols discussed in Sec. IX.C). It is noteworthy, however, that, as mentioned in Sec. III.C.3, by explicitly limiting the adversary's capabilities in a composable framework, we obtain a security definition that is equivalent to weak security criteria that have appeared in the literature, e.g., accessible information.

VII. SECURE CLASSICAL MESSAGE TRANSMISSION

One of the main tasks in cryptography is to securely transmit a confidential message from one player to another. Securely transmitting the message means that the adversary does not learn anything about the message (except unavoidable leaks such as the message length) and cannot modify the message either. We also want to achieve this with minimal assumptions on the available resources. In Fig. 3 we depict the steps necessary to construct such a secure channel from nothing but insecure channels and an initial short key. The aim of this section is to explain this construction in detail.

In Sec. VII.A we first show how to construct an authentic channel that is used by both QKD and the OTP. In Secs. VII.B and VII.C we then revisit the notions of a secure key and secure channel resources introduced earlier and discuss a modification used here.[38] In Sec. VII.D we put the individual parts of the construction together and show how this gives a construction of a secure channel from a short secret key and insecure communication channels only.

A. Authentication

The use of an authentic channel is essential for many cryptographic protocols, including quantum key distribution, as we saw earlier. It allows the players Alice and Bob to be sure that they are communicating with each other and not with an adversary Eve. The authentic channel we used in Sec. III (such as Fig. 7) is, however, idealized: it guarantees that the recipient always receives the message that was sent. In a realistic situation, one has to assume that an adversary may jumble or cut the communication and prevent messages from arriving. What still can be constructed is a channel that guarantees that Bob does not receive a corrupted message. He receives either the correct message sent by Alice or an error that indicates an attempt by Eve to change the message. This can be modeled by giving Eve's idealized interface two controls: the first provides her with Alice's message, and the second allows her to input 1 bit that specifies whether Alice's

[38]We take into account the possibility that Eve may prevent the honest players from obtaining the key or the transmitted message, which was ignored earlier for simplicity.

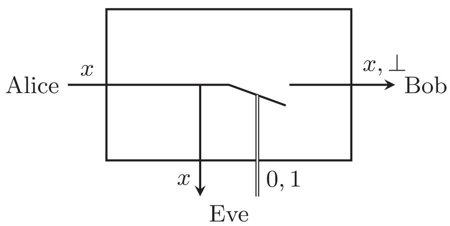

FIG. 25. An authentic channel resource. The message input at Alice's interface is visible to Eve, who gets to decide whether Bob receives it or not. But this guarantees that if Bob does receive a message, it corresponds to the one sent by Alice.

message should be delivered or Bob gets an error instead. We illustrate this in Fig. 25.

As later explained, such an authentic channel can be constructed from a completely insecure channel together with a shared secret key. Although this may be done using a nonuniform secret key [see Renner and Wolf (2003), Dodis and Wichs (2009), and Aggarwal *et al.* (2019)], we review here a simpler construction, originally proposed by Wegman and Carter (1981), that still needs only a short key, but that, however, has to be close to uniform: one computes a hash $h_k(x)$ of the message x and sends the string $x\|h_k(x)$ to Bob, where k is the short shared secret key and $\{h_k\}_{k\in\mathcal{K}}$ is a family of strongly universal hash functions.[39] Alice's part of the authentication protocol π_A^{auth} thus gets as input a key k from an ideal key resource, as well as a message x from Alice, and sends $x\|h_k(x)$ over the insecure channel. When Bob receives a string $x'\|y'$, he needs to check whether $y' = h_k(x')$. His part of the protocol hence gets as input the key k from the ideal key resource, the message $x'\|y'$ delivered by the channel, and output x' if $y' = h_k(x')$ or, otherwise, a symbol $\perp$ to indicate an error. This is depicted in Fig. 26.

To capture completeness of this protocol, one considers instead of the insecure channel $\mathcal{C}$ as in Fig. 26 a noiseless channel with a blank interface for Eve[40] [as illustrated in Fig. 1(a)], and the constructed channel is also a perfect noiseless channel instead of the channel from Fig. 25. These real and ideal systems are indistinguishable, as they are both identity channels that faithfully transmit x from Alice to Bob. This proves completeness, and we can therefore focus in the following on the other part, namely, proving the soundness of the protocol.

In the ideal setting, the authentic channel (Fig. 25) has the same interface on Alice's and Bob's sides as the real setting (Fig. 26): Alice can input a message, and Bob receives either a message or an error. However, Eve's interface looks quite different: in the real setting she can modify the transmission on the insecure channel, whereas in the ideal setting the adversarial interface provides only controls to read the

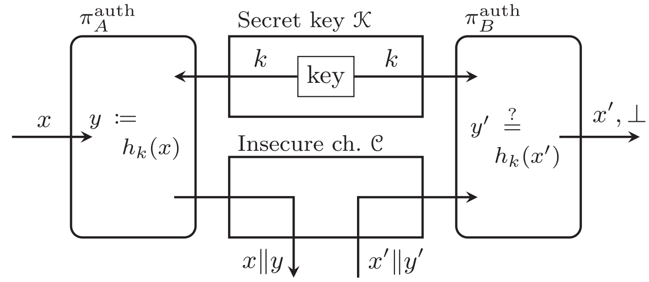

FIG. 26. The real authentication system consists of the authentication protocol $(\pi_A^{\mathrm{auth}}, \pi_B^{\mathrm{auth}})$ as well as the secret key and insecure channel resources $\mathcal{K}$ and $\mathcal{C}$. As in previous illustrations, Alice has access to the left interface, Bob has access to the right interface, and Eve has access to the lower interface.

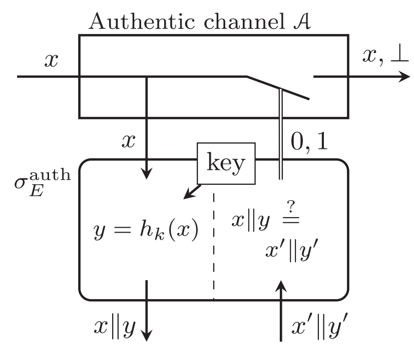

FIG. 27. The ideal authentication system (Alice has access to the left interface, Bob has access to the right interface, and Eve has access to the lower interface) consists of the ideal authentication resource and a simulator σ_E^{auth}.

message and interrupt the transmission. From Definition 1 we know that an authentication protocol constructs the authentic channel if there is a simulator σ_E^{auth} that can recreate the real interface while accessing only the idealized one. A choice for the simulator is to first generate its own key k and output $x\|h_k(x)$. Upon receiving $x'\|y'$, it then checks if $x'\|y' = x\|h_k(x)$ and presses the switch on the authentic channel to output an error if this does not hold. We illustrate this in Fig. 27.

In this case, an authentication protocol is ε secure if Figs. 26 and 27 are ε close, i.e.,

$$\pi_A^{\mathrm{auth}}\pi_B^{\mathrm{auth}}(\mathcal{K}\|\mathcal{C}) \approx_\varepsilon \mathcal{A}\sigma_E^{\mathrm{auth}}.$$

Original works defining authentication (Wegman and Carter, 1981; Simmons, 1985, 1988; Stinson, 1990, 1994) did not use such a composable security definition. Instead, they considered two kinds of attacks. In the first, the adversary obtains a pairing of a valid message and an authentication tag and tries to find a pairing of a different message and corresponding valid authentication tag: this is called a substitution attack. In the second, the adversary directly tries to find a pair of message and corresponding valid authentication tag: this is called an impersonation attack. It was then shown that if the family of hash functions used are ε-almost strongly

[39]See the formal definition later in this work in footnote 41.

[40]Unlike in the case of QKD, we do not consider noisy channels between Alice and Bob, as such noise could be removed easily by encoding the communication with an appropriate classical error correcting code. Therefore, the assumed and constructed channels faithfully deliver the message from Alice to Bob.

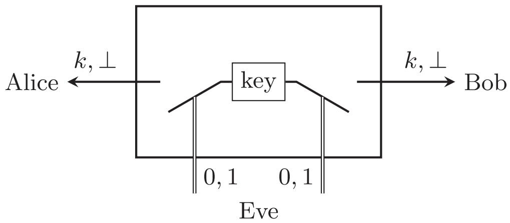

FIG. 28. A secret key resource allowing Eve to control who gets the key: the2 bits Eve inputs control whether Alice or Bob obtain, respectively, a key or an error message from the resource.

two universal,[41] the probability of either of these attacks being successful is bounded by ε. A composable security proof for these schemes was given by Portmann (2014), who showed that Eq. (36) is satisfied, again under the condition that the family of hash functions used are ε-almost strongly two universal.[42]

Note that a distinguisher interacting with either the real or the ideal system has a choice between providing messages in two different orders. It can first provide Alice with a message, receive the ciphertext[43] at Eve's interface, then input a modified ciphertext, and finally learn whether the ciphertext is accepted or not at Bob's interface. Or it can first input a forged ciphertext at Eve's interface, then learn if it is accepted, and finally provide Alice with a message and obtain the corresponding ciphertext at Eve's interface. These two orders of messages roughly correspond to the substitution and impersonation attacks.

The secret key resource used thus far in this section assumes that both players always get a copy of the key. However, in Sec. III we modeled a secret key resource with a switch at Eve's interface, giving her the possibility to prevent the players from getting the key. If such a switch is present and Eve flips it, the honest players will not be able to run the authentication protocol at all. This does not, however, change the ideal resource constructed, because not running the protocol or running the protocol but with Eve preventing the message from being delivered are essentially equivalent. However, the proof in this case requires a different simulator, one that receives the bit deciding whether the players get a key or not and then acts accordingly.

In Sec. VII.B an even weaker secret key resource is considered, one that allows Eve to decide if only one of the two players gets a secret key and not the other; this is drawn in Fig. 28. With reasoning similar to that used previously, one can see that this does not change the outcome of the protocol either: it still constructs the authentic channel from Fig. 25.

B. Quantum key distribution

In Sec. III we analyzed QKD protocols that use an insecure quantum channel and an authentic channel with the guarantee that the message is always delivered, as indicated in Fig. 7. The motivation behind this standard choice was that if the message is not delivered, then the players abort and the scheme is trivially secure. In other words, the nontrivial case that needs to be analyzed to prove that a QKD scheme constructs the ideal key resource is the one in which the adversary does not use her switch and allows the messages to be delivered on the authentic channel.

Nonetheless, if we do replace the authentic channel with the one that can actually be constructed from Fig. 25, then we also have to weaken the ideal key resource that is constructed. If the players get an error message from the authentic channel instead of the intended message, they will simply abort the protocol and not produce a key. In the version from Sec. III, Eve already has the power to prevent the players from getting a secret key. The difference is that now Eve can let one player get the secret key but not the other, such as by jumbling the last message between the players. The resulting ideal key resource is drawn in Fig. 28.

The analysis carried out in Sec. III goes through with only minor changes with the weaker authentic channel and secret key resources since the only differences are the abort conditions, which now may additionally occur because of failed authentication. In particular, the reduction from the constructive statement [Eq. (8)] to the trace distance criterion [Eq. (14)] is unaffected by these changes of resources.

Note that if Eve prevents one player from getting the key but not the other, the players are generally unaware of this fact and may end up with mismatching key lengths.[44] This is a problem that is not specific to QKD but rather happens in general with any key distribution scheme (Wolf, 1999).

C. One-time pad

In Sec. II.C we analyzed the OTP and showed that it constructs a secure channel given a secret key and an authentic channel. But again, we used an authentic channel that guarantees that the message is transmitted, as well as a secret key that is guaranteed to be delivered to the players. It is easy, however, to convince oneself that these extra assumptions about the resources do not affect the security of the protocol, since if the players do not get a key or a message they simply abort the protocol, in which case security holds trivially. Therefore, if we plug in the authentic channel from Fig. 25 and

[41]A family of functions $\{h_k : \mathcal{X} \to \mathcal{Y}\}_k$ is said to be ε-almost strongly two universal if any two different messages are almost uniformly mapped to all pairs of tags, i.e., $\forall x_1, x_2, y_1, y_2, x_1 \neq x_2, \Pr_k[h_k(x_1) = y_1 \text{ and } h_k(x_2) = y_2] \leq \varepsilon/|\mathcal{Y}|$ (Stinson, 1994). The family of functions is said to be strongly two universal if $\varepsilon = 1/|\mathcal{Y}|$.

[42]Portmann (2014) additionally showed that part of the secret key k can be recycled since only a number of bits corresponding to the length of the hash $h_k(x)$ are leaked. This is discussed in Sec. VIII.B.

[43]By ciphertext we denote the pairing of the message and authentication tag generated by the sender.

[44]The protocol may be designed in such a way that the round in which Eve needs to jumble the communication so that one player accepts the key but not the other is unknown to her. Thus, her probability of success will be $O(1/n)$, where n is the number of rounds of communication.

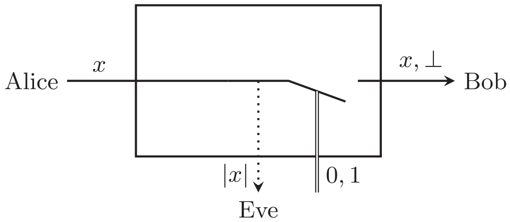

FIG. 29. A secure channel that allows Eve to learn the length of the message and prevent Bob from receiving it. When Alice inputs a message x at her interface, information about the length of the message is given to Eve, who can additionally press a switch that either delivers Alice's message to Bob or instead provides him with an error message $\perp$.

the secret key resource from Fig. 28, we merely need to weaken the secure channel that is constructed so that it may abort as well. We can model this by adding a switch at Eve's interface that, when flipped, delivers an error message at Bob's interface instead of Alice's message, as depicted in Fig. 29. This is similar to the authentic channel resource from Fig. 25, except that Eve receives only the length of the message rather than the message itself.

The analysis of the OTP with these altered resources is identical to the one in Sec. II.C since the only difference is that both the real and ideal systems might abort or output an error at Bob's interface instead of a message if Eve operates the corresponding switch at her interface.

D. Combining the subprotocols

Let $\mathcal{A}$ denote an authentic channel resource, as illustrated in Fig. 25, and let $\mathcal{K}^\ell$ denote a secret key resource of length ℓ, as drawn in Fig. 28. Furthermore, let $\mathcal{C}$ be an insecure classical, and let $\mathcal{Q}$ be an insecure quantum channel. Finally, we denote by $\mathcal{S}$ a secure classical channel, as depicted in Fig. 29. If we summarize the results presented thus far, we see that an authentication protocol constructs an authentic channel from an insecure channel and a secret key, i.e.,

$$\mathcal{K}^a \| \mathcal{C} \xrightarrow{\pi_{AB}^{\text{auth}}, \varepsilon_{\text{auth}}} \mathcal{A}, \tag{36}$$

a QKD protocol constructs a shared secret key resource from an authentic channel and an insecure quantum channel, i.e.,

$$\mathcal{A} \| \mathcal{Q} \xrightarrow{\pi_{AB}^{\text{QKD}}, \varepsilon_{\text{QKD}}} \mathcal{K}^n,$$

and a OTP constructs a secure channel from an authentic channel and a secret key, i.e.,

$$\mathcal{A} \| \mathcal{K}^m \xrightarrow{\pi_{AB}^{\text{OTP}}, 0} \mathcal{S}.$$

Using also the fact that a key can be split, i.e.,

$$\mathcal{K}^{a+b} \xrightarrow{id, 0} \mathcal{K}^a \| \mathcal{K}^b,$$

denoting by a_{QKD} the length of the key used by the authentication subroutines for QKD[45] and by a_{OTP} the length of the key used to authenticate the message for constructing the secure channel, we obtain

$$\mathcal{C} \| \mathcal{C} \| \mathcal{Q} \| \mathcal{K}^{a_{\text{QKD}}} \xrightarrow{\pi_{AB}, \varepsilon} \mathcal{S} \| \mathcal{K}^{n-m-a_{\text{OTP}}},$$

where π_{AB} is the composition of all the protocols and $\varepsilon = \varepsilon_{\text{QKD}} + \varepsilon_{\text{auth}}^{\text{QKD}} + \varepsilon_{\text{auth}}^{\text{OTP}}$ is the sum of the errors of the individual protocols. We depict this in Fig. 30, where for simplicity we have drawn only one round of authentication as a subroutine of QKD.

VIII. OTHER CRYPTOGRAPHIC TASKS

In our description thus far we have adopted the composable view and regarded cryptographic protocols as constructions; i.e., protocols construct some resources given other resources.[46] Cryptographic protocols proposed in the literature have not always been defined in this way, but instead are specified in terms of particular securitylike properties, e.g., that an adversary is unable to guess the content of an encrypted message. Sometimes these properties can be rephrased as an ideal system within the real-world ideal-world paradigm, as discussed in Sec. III.C.3. In this section we review some of the major results in quantum cryptography from this perspective; i.e., we present them as constructive statements, defining the resources constructed and used by the protocols. For a broader review of quantum cryptography, see the recent survey by Broadbent and Schaffner (2016).

A. Secure quantum message transmission

From a theory of resources perspective, the task of securely transmitting a quantum message from Alice to Bob is nearly identical to the corresponding classical task, which is analyzed in Sec. VII. Here too we require the players to share a secret key resource and an insecure channel, and the goal is to construct a secure channel, with the only difference being that the insecure and secure channels are both quantum channels. We already encountered insecure quantum channels in Sec. III, where they were used for QKD. A secure quantum channel is modeled analogously to a secure classical channel as drawn in Fig. 29, except that the messages sent are quantum. We depict this in Fig. 31.

The first protocols that construct such a secure quantum channel from a shared secret key and an insecure channel were proposed by Barnum *et al.* (2002). They follow the same pattern as classical message transmission: one first encrypts the quantum message with a quantum OTP, then encodes it in a larger space so as to detect any errors that may be introduced by an adversary. However, contrary to the case of classical

[45] A QKD protocol usually authenticates many messages, and they may be going in both directions between Alice and Bob. For simplicity, we write this here as just one round of authentication, which uses a key of length a_{QKD} and has error $\varepsilon_{\text{auth}}^{\text{QKD}}$.

[46] See also footnote 9 for other uses of resource theories in quantum mechanics.

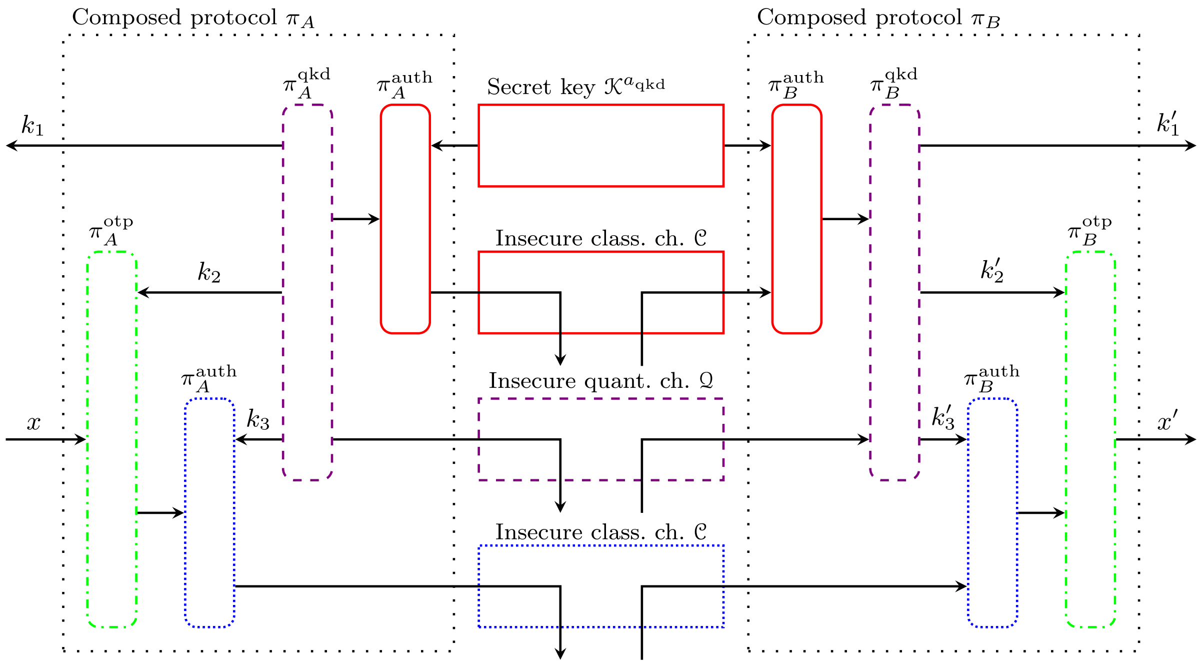

FIG. 30. Composition of QKD, authentication, and OTP protocols. For simplicity, we have drawn only one round of authentication as a subroutine of QKD as $\pi_{AB}^{\text{auth}}(\mathcal{K}^{a_{\text{QKD}}}\|\mathcal{C})$ (red solid line). The QKD protocol π_{AB}^{QKD} (violet dashed line) constructs a shared key resource that produces the long key (k_1, k_2, k_3). The second authentication protocol π_{AB}^{auth} (blue densely dotted line) then uses part of this key to construct another authentic channel, and the OTP protocol π_{AB}^{OTP} (green dash-dotted line) uses another part of this key to encrypt and decrypt the message sent on the channel.

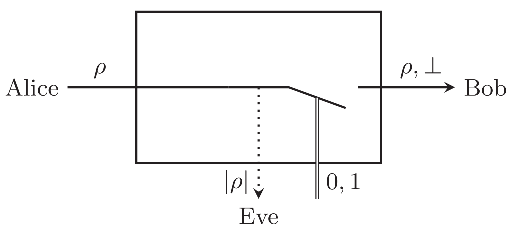

FIG. 31. A secure quantum channel that allows Eve to learn the length of the quantum message (informally denoted as $|\rho|$) and prevent Bob from receiving it. When Alice inputs a message ρ at her interface, information about the length of the message is given to Eve, who can additionally flip a switch that either delivers Alice's message to Bob or provides him with an error message $\perp$.

messages, there is no known way to view these two steps as two distinctive constructive statements, i.e., as a construction of an authentic channel from an insecure channel and a second construction of a secure channel from an authentic channel. This means that the analysis has to include both aspects at the same time.

In Sec. VIII.A.1 we explain this construction and in Sec. VIII.A.2 we review additional work on the topic. At the end of this section (in Sec. VIII.F) we revisit this topic from a computational perspective.

1. Generic protocol

The classical OTP introduced in Sec. II.C can be seen as randomly flipping each bit of the message. The quantum OTP (Ambainis *et al.*, 2000; Boykin and Roychowdhury, 2003) follows the same principle: one flips the bits and phases of the message at random. For $x, z \in \{0,1\}^n$, let X^x and Z^z denote operators on $(\mathbb{C}^2)^{\otimes n}$ that perform bit and phase flips in positions indicated by the strings x and z, respectively. The quantum OTP consists of choosing x and z uniformly at random and applying the corresponding operation to the message. For any state ρ_{MR}, where M is a register of size 2^n, we thus have

$$\frac{1}{2^{2n}}\sum_{x,z} Z^z X^x \rho_{MR} X^x Z^z = \tau_M \otimes \rho_R,$$

where τ_M is the fully mixed state and ρ_R is the reduced density operator of ρ_{MR}. We call the operators $Z^z X^x$ defined in this way Pauli operators.[47]

The second ingredient needed to construct a secure quantum channel is an error correcting code that is going to be used to detect errors in the transmission, i.e., tampering by an adversary. Generally, an error correcting code may be seen as a map from a message space $\mathcal{H}_M$ to a larger physical space $\mathcal{H}_C$. For simplicity, we model the encoding for a code $\mathcal{C}_k$ as first appending a state $|0\rangle \in \mathcal{H}_T$ to the message ρ_M, where $\mathcal{H}_C = \mathcal{H}_M \otimes \mathcal{H}_T$, followed by applying a unitary U_k to the resulting state, e.g., $\sigma_C = U_k(\rho_M \otimes |0\rangle\langle 0|)U^\dagger{}_k$.

To test to see whether an error occurred, one decodes the received state $\tilde{\sigma}_C$ by applying the inverse operation $U^\dagger{}_k$ and measures the T register in the computational basis. If the result

[47]This notation simplifies the presentation here but deviates from the more commonly used definition of the Pauli operators as $i^{z\cdot x}Z^z X^x$, where $z \cdot x = \sum_j z_j x_j$ and i is the imaginary unit.

is not 0, this is evidence of noise. We say that a code detects an error V if, after decoding a message to which this error was applied (i.e., $\tilde{\sigma}_C = V\sigma_C V^\dagger$), one always gets a measurement outcome other than 0. Furthermore, we call an error trivial if it never affects the code word, i.e., for any ρ_M,

$$U^\dagger{}_k V U_k(\rho_M \otimes |0\rangle\langle 0|)U^\dagger{}_k V^\dagger U_k = \rho_M \otimes |0\rangle\langle 0|.$$

For an error V to modify a message and yet not be caught, it must be nontrivial and not detected by the code used. For the purpose of constructing a secure channel according to the method of Barnum *et al.* (2002), it is sufficient to detect all Pauli errors. Barnum *et al.* (2002) defined a set of codes that they called purity testing codes. They guarantee that with high probability over the choice of code all Pauli errors are either caught or trivial. More precisely, a set $\{\mathcal{C}_k\}_k$ of codes forms a family of ε-purity testing codes if, for any Pauli error $X^x Z^z$, the probability over a uniformly random choice of k that this error is neither caught nor trivial is less than ε.

The protocol for constructing a secure channel then works as follows. The sender first encrypts the message with a quantum OTP, i.e., a Pauli operator $Z^z X^x$ chosen uniformly at random according to the secret key. A state $|s\rangle$ is then appended to the message, where s is also chosen uniformly at random according to the key. Finally, the resulting state is encoded with a unitary U_k corresponding to the encoding operation of an element of a purity testing code family $\{\mathcal{C}_k\}_k$, where again k is chosen uniformly at random according to the secret key. Decryption works in the following way: the receiver applies the inverse operator $U^\dagger{}_k$ and then measures the T register. If the outcome is not s, the message was jumbled and the player outputs an error symbol. Otherwise, the receiver decrypts the message with the operator $Z^z X^x$.

2. Concrete schemes

Barnum *et al.* (2002) introduced the general family of protocols described in Sec. VIII.A.1, and also provided a concrete construction of a purity testing code family that has good parameters. Following this seminal work, a variety of further protocols for authentication of quantum messages based on different codes have been proposed in the literature. Authentication using the signed polynomial code (Ben-Or *et al.*, 2006; Aharonov, Ben-Or, and Eban, 2010), the trap code (Broadbent, Gutoski, and Stebila, 2013; Broadbent and Wainewright, 2016), the Clifford code (Aharonov, Ben-Or, and Eban, 2010; Dupuis, Nielsen, and Salvail, 2012; Broadbent and Wainewright, 2016) [which is a unitary 3-design (Webb, 2016; Zhu, 2017)], a unitary 8-design (Garg, Yuen, and Zhandry, 2017), and a unitary 2-design[48] (Alagic and Majenz, 2017; Portmann, 2017a) are all instances of the family from Barnum *et al.* (2002), with alternative purity testing codes.[49] To the best of our knowledge, only the Auth-QFT-Auth scheme from Garg, Yuen, and Zhandry (2017) is not known to follow the model of Barnum *et al.* (2002).

Although most of these works provide some kind of security proof for the protocol, only two papers consider a composable security definition, namely, those of Hayden, Leung, and Mayers (2011) and Portmann (2017a). Both works show that the family of protocols from Barnum *et al.* (2002) construct a secure quantum channel from a shared secret key and an insecure quantum channel. Note, however, that Hayden, Leung, and Mayers (2011) considered a restricted class of distinguishers [those that perform a substitution attack (see Sec. VII.A)] and Portmann (2017a) analyzed only a subset of this family, for which the purity testing code family detects all (rather than only the nontrivial) errors with high probability.[50] In fact, both papers prove that one may additionally recycle part of the key, as we discuss further in the following section.

B. Key reuse in classical and quantum message transmission

As mentioned in Secs. VII.A and VIII.A.2, part of the key used in the constructions of secure channels may be recycled; i.e., at the end of the protocol it can be added back to a pool of secret key bits. For example, in the case of the classical message authentication analyzed in Sec. VII.A, the sender appends a tag $y = h_k(x)$ to the message x. The value of the tag y depends on the shared secret key k, and every bit of the tag leaks (at most) a bit of the secret key to the adversary. But the key is longer than the tag, so the bits that are not leaked may be reused. It is vital, however, that they not be recycled too soon. If the sender reuses part of the key before the receiver obtains the authenticated message, the adversary may learn these bits and use this information to successfully change the message and authentication tag.

To recycle key bits in classical or quantum message transmission, the real system is changed as follows. The players first need an extra resource, a 1-bit backward authentic channel, allowing the receiver to tell the sender whether the message was successfully received. Once this confirmation is sent, the receiver may recycle part of the key; i.e., it is output by the corresponding converter. And once this confirmation is received by the sender, she also recycles the same part of the key. The ideal resource constructed in this way corresponds to the parallel composition of a secure (or authentic) channel and a secret key resource (drawn in Fig. 32). As previously, the adversary may control whether the message is delivered on the secure channel. And since the amount of the key that is recycled may depend on the adversary's behavior as well; namely, by allowing or preventing the message from being delivered, we model the resource as being equipped with a switch to control how much of the key the players get.

In the case of authentication of classical messages, Wegman and Carter (1981) proposed that part of the key can be safely recycled. Here, if the two-universal hash function has the special form $h_{k_1,k_2}(x) = f_{k_1}(x) \oplus k_2$, where k_2 is a bit string

[48]Because any unitary t-design for $t \geq 2$ is a unitary 2-design, it follows that any t-design constructs a secure quantum channel.

[49]Most of these works consider a construction where the message is first encoded with a purity testing code and then encrypted. But, as shown in Portmann (2017a), this is equivalent to the original scheme of Barnum *et al.* (2002), which reverses the order of these two operations.

[50]One refers to this as a strong purity testing code.

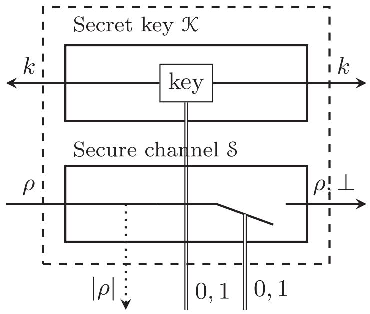

FIG. 32. The ideal system for a secure channel with key recycling. It consists of a secure channel $\mathcal{S}$ and key resource $\mathcal{K}$. The adversary controls the length of the recycled key (through her input to $\mathcal{K}$), as well as whether the receiver obtains the message (through her input to $\mathcal{S}$).

of the same length as the tag, then k_1 may be recycled, but a new k_2 is needed for every message. It was proven by Portmann (2014) that this scheme is composable and constructs the previously described ideal resource.

In the case of quantum messages, roughly the same holds in the case where the message fails the authentication, The number of bits of key leaked depends on the length of the ciphertext, and the rest can be recycled.[51] But in the case where the message is accepted, the players can recycle more of the key. This holds because of the no-cloning principle of quantum mechanics: if the receiver holds the correct ciphertext, then the adversary cannot have a copy of it and thus does not hold any information about the key either. It was first shown by Hayden, Leung, and Mayers (2011) that nearly all of the key could be recycled in the case in which the message is accepted. Portmann (2017a) then showed that every bit of the key can indeed be recycled. This is not known to hold for all schemes that construct a secure quantum channel, but so far only for those that use strong purity testing codes (Portmann, 2017a).

C. Delegated quantum computation

The setting in which a client, typically with bounded computational resources, asks a server to perform some computation for her is called delegated computation. The client might not want the server to learn what computation it is performing for her and might want to run a protocol that hides the underlying computation; this property is called blindness in the literature. Furthermore, the client might want to verify that the server correctly performed the computation she required; this is known as verifiability.

The task of delegating a quantum computation was first studied by Childs (2005), with the goal of achieving blindness. In follow-up works, the requirements on the client's information-processing abilities were reduced. Broadbent, Fitzsimons, and Kashefi (2009) proposed the first protocol for blind delegated quantum computation that does not require the client to have quantum memory, but instead only the

[51]If the ciphertext is n qubits long, about $2n$ bits of key are lost; see Portmann (2017a) for the exact parameters.

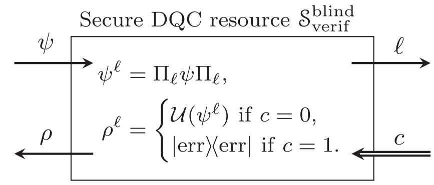

FIG. 33. Ideal DQC resource. The client has access to the left interface, and the server has access to the right interface. The server obtains some information ℓ about the input and can decide whether the client gets the correct outcome or an error by inputting a bit c.

ability to prepare different pure states. This result was extended by Fitzsimons and Kashefi (2017) to include verifiability as well.

DQC was formalized as a constructive statement by Dunjko *et al.* (2014). They modeled a DQC protocol that achieves both blindness and verifiability as constructing a resource $\mathcal{S}^{\mathrm{blind}}_{\mathrm{verif}}$ that works as follows. It first receives a description of the required computation as a state ψ from the client. Every computation necessarily leaks some information to the server, such as an upper bound on the computation size, so the resource computes this leaked information ℓ and outputs it at the server's interface. The server can then decide if it will cheat (in which case the client will, however, get an error message) or output the correct result of the computation, which is evaluated by applying an operator $\mathcal{U}$ to the input. This is depicted in Fig. 33. A DQC protocol constructs such a resource from nothing more than a shared communication channel between client and server.

A weaker resource that provides only blindness and not verifiability can be obtained by increasing the power of the server at its interface of the resource. Instead of inputting a bit that decides whether the client gets the correct outcome, the server can decide what output the client gets, but still receives only the leaked information ℓ (Dunjko *et al.*, 2014).

Dunjko *et al.* (2014) showed that the protocols from Broadbent, Fitzsimons, and Kashefi (2009) and Fitzsimons and Kashefi (2017) satisfy the corresponding constructive definitions. These protocols still require the client to prepare a few different single qubit quantum states and send them to the server. To better analyze this requirement, Dunjko and Kashefi (2016) decomposed the construction of a DQC resource into two steps. They first consider a resource that provides the server with the random states that it needs (and the client with a description of which state was given to the server); then the DQC protocol constructs the DQC resource given this state preparation resource. This decomposition then allowed Gheorghiu and Vidick (2019) to design a protocol that constructs the required state preparation resource for an entirely classical client: this was achieved by sacrificing information-theoretic security for computational security. Composing this with a DQC protocol, one gets DQC for an entirely classical client, albeit with computational security. This is believed not to be possible with information-theoretic security (Aaronson *et al.*, 2019). We note, however, that Gheorghiu and Vidick (2019) made a nonstandard assumption about available resources, without which such a result does not seem possible (Badertscher *et al.*, 2020).

It is instructive to compare this to early definitions of blindness, such as those from Broadbent, Fitzsimons, and Kashefi (2009) and Fitzsimons and Kashefi (2017). In those definitions the requirement is that the server learns nothing about the computation except for the allowed leaked information ℓ. Roughly speaking, this means that the state ρ^{ψ^ℓ} held by the server at the end of the protocol, where ψ^ℓ is an input that leaks information ℓ, must be such that

$$\rho^{\psi^\ell} \approx \rho^\ell. \tag{37}$$

In other words, it depends only on ℓ, not on any other part of the input. If we compare this to the constructive definition of Dunjko *et al.* (2014), in which the distinguisher has access to both the server's interface and the client's interface of the resources, Eq. (37) corresponds to the special case where maximize not over all distinguishers, but rather over only those that ignore the output received by the client. Following Sec. III.C.3 one may express this as a restriction on the resource constructed instead of a restriction on the distinguisher. Requiring a DQC protocol to satisfy Eq. (37) is mathematically equivalent to requiring it to construct an ideal resource that does not provide the client with the result of the computation.

D. Multiparty computation

In this section we consider a setting where multiple mutually distrustful parties wish to evaluate a possibly randomized function to which each of them provides an input, or they wish to jointly evaluate a CPTP map on shared quantum inputs. They, however, do not want the other parties to learn anything about their input other than what can be learned from the output. Furthermore, they also want to guarantee that if they get an output, then this is the correct output. An example is a function that outputs which player i has the largest input x_i; e.g., the players want to know who earns more without revealing their salary to the others. Another example is generating a random coin flip, in which case no input is required. Generally speaking, multiparty computation corresponds to constructing an ideal resource that first takes the inputs from all parties and then provides them with the correct output.

1. Bit commitment

A bit commitment resource is a system in a two-party setting that forces a player (say, Alice) to commit to a value without revealing this value to the other player (say, Bob). At a later point, the commitment is "opened" so that Bob may know which value Alice committed to. More precisely, Alice sends a bit b to the resource, and Bob is notified that Alice is committed to a value. Alice may then send an open command to the resource, at which point b is delivered to Bob. In the classical setting, such a resource cannot be constructed from communication channels alone (Canetti and Fischlin, 2001; Maurer and Renner, 2011), but extra resources such as a common reference string (a random string shared by all parties) are needed (Canetti and Fischlin, 2001).

The argument of Maurer and Renner (2011) was extended by Vilasini, Portmann, and Rio (2019) to prove that even if the players use quantum protocols and even if the adversary is computationally bounded, has only bounded or noisy storage, and is restricted by relativistic constraints,[52] it is still impossible to construct a bit commitment resource without further setup assumptions than communication channels.

It has been suggested that one could construct bit commitment if one takes relativity into account, i.e., that messages cannot be sent faster than the speed of light (Kent, 1999, 2012; Kaniewski *et al.*, 2013). However, these protocols do not construct a bit commitment resource: Appendix A of Kaniewski (2015) [][53] proved that if one composes these relativistic bit commitment protocols with the protocol from Unruh (2010) to construct oblivious transfer from bit commitment,[54] then the result is not a secure oblivious transfer protocol. It has now been proven that taking relativity into account is not sufficient to achieve bit commitment (Vilasini, Portmann, and Rio, 2019), which we discuss in more detail in Sec. VIII.E.

2. Coin flipping

Another well studied resource is that of coin flipping, in which a random coin is flipped and both players are provided with the result. The impossibility proof for bit commitment of Maurer and Renner (2011) can be adapted to show that coin flipping and biased coin flipping (where a player is allowed to partially bias the flip) are also impossible without further assumptions. Note that this proof is valid independently of whether one considers classical or quantum strategies. A direct proof for the impossibility of coin flipping in the quantum and relativistic setting (even in the case of computational- and memory-bounded adversaries) was given by Vilasini, Portmann, and Rio (2019).

Coin expansion is a weaker task in which one constructs a resource that produces a sequence of coin flips from a weaker resource that produces fewer coin flips. This has been shown to be impossible for classical players with information security (Hofheinz, Müller-Quade, and Unruh, 2006; Seiler and Maurer, 2016) but is possible with computational security (Hofheinz, Müller-Quade, and Unruh, 2006) and remains open in the quantum case.

3. Two-party function evaluation and oblivious transfer

It was shown by Ishai, Prabhakaran, and Sahai (2008) that a resource that evaluates any classical probabilistic polynomial-time function with two inputs can be constructed given an oblivious transfer resource, i.e., a system that receives two

[52]This means that the adversary cannot send information between two points faster than it takes light to travel between the two points; see Sec. VIII.E.

[53]The proof in Appendix A of Kaniewski (2015) used the same attack as Brassard *et al.* (1998), who showed that some noncomposable definitions of bit commitment appearing in the classical literature cannot be used to force a quantum player to commit to a measurement outcome.

[54]Oblivious transfer and the construction from Unruh (2010) are discussed in Sec. VIII.D.3.

strings s_0 and s_1 from one player Alice and a bit c from the second player Bob and sends Bob s_c.

In the quantum setting, it is possible to construct an oblivious transfer resource from a bit commitment resource. The construction of oblivious transfer from bit commitment was first proposed by Crépeau and Kilian (1988), was adapted to noisy channels by Bennett, Brassard *et al.* (1992), and was proven secure by Unruh (2010). When this result is combined with that of Ishai, Prabhakaran, and Sahai (2008), it follows that bit commitment is universal for classical two-party computation (Unruh, 2010).

It is not possible, however, to construct an oblivious transfer resource from nothing but communication channels, even if the adversary is computationally bounded, has only bounded or noisy storage, and is restricted by relativistic constraints (Laneve and Rio, 2021).

4. Everlasting security

Unruh (2013) studied multiparty computation in the setting of everlasting security. This means that one relies upon a computational assumption, but this assumption has to be broken during the execution of the protocol for an adversary to break the scheme. If this is not the case, then even a computationally unbounded adversary cannot achieve a significant advantage after the protocol has terminated. This is generally not satisfied by computational encryption schemes, because an adversary could obtain a ciphertext and wait for an advancement in algorithms to break the scheme and obtain the message. But if a computational authentication scheme is executed, then the adversary must be able to perform the hard computation before the message is received and authenticated.

Composition in such a setting is not straightforward, and Unruh (2013) provided the necessary definition that a scheme must satisfy to be composable. He showed how to perform authentication given a signature card that when composed with QKD and secure encryption as in Sec. VII, results in a secure channel. He also provided a way to perform bit commitment based on signature cards. Composing this with the protocols from Sec. VIII.D.3 allows one to perform any multiparty computation with everlasting security.

5. Multiparty quantum computation

The tasks studied thus far in this section are concerned with multiparty evaluation of a classical function [so-called multiparty computation (MPC)], but while using quantum communication and computation to possibly achieve what cannot be done classically. The problem of multiparty quantum computation (MPQC) generalizes this to the case in which the inputs and outputs are quantum (Crépeau, Gottesman, and Smith, 2002; Ben-Or *et al.*, 2006; Dupuis, Nielsen, and Salvail, 2012; Dulek *et al.*, 2020; Lipinska, Ribeiro, and Wehner, 2020; Alon *et al.*, 2021). The relation between inputs and outputs is then most generally described by a CPTP. It is standard to use a composable framework for analyzing classical MPC (Cramer, Damgård, and Nielsen, 2015). But to the best of our knowledge the only work on MPQC that mentions that the results hold in a composable framework is that of Ben-Or *et al.* (2006), and they provided only a proof sketch. All other works assumed that the dishonest party performs their attack in an isolated way, interacting with the environment (the distinguisher) only before the protocol starts and after the protocol ends. This is the so-called stand-alone security model, and protocols proven secure in such a model do not necessarily compose concurrently with other protocols. In particular, they might not be secure if two instances of the same protocol are run in parallel. Nonetheless, for MPQC we do not know of any attacks on protocols that have been run concurrently, and it is plausible that exactly the same result would go through in a composable security framework.

The ideal resource that one wants to construct in MPQC receives the inputs from all parties, performs the quantum computation, and then provides each player with their part of the output. Crépeau, Gottesman, and Smith (2002), Ben-Or *et al.* (2006), and Lipinska, Ribeiro, and Wehner (2020) considered an ideal resource that is guaranteed to provide the output to the honest players. Crépeau, Gottesman, and Smith (2002) first showed that this can be achieved if the fraction of dishonest parties is $t < n/6$, where n is the total number of players. Ben-Or *et al.* (2006) improved this to $t < n/2$ cheating parties. Lipinska, Ribeiro, and Wehner (2020) decreased the number of qubits and communication complexity needed to get the same result.

Dupuis, Nielsen, and Salvail (2012), Dulek *et al.* (2020), and Alon *et al.* (2021) defined the ideal resource such that it first provided the dishonest parties with their share of the output. They then provided a bit to the ideal resource, which indicates whether the honest parties should receive their output or an abort symbol instead. This is called unfairness. Weakening the ideal resource in this way allows the number of dishonest parties to be any $t < n$. Dupuis, Nielsen, and Salvail (2012) first showed how to do this in the two-party case. Dulek *et al.* (2020) extended this to the multiparty setting. Alon *et al.* (2021) improved the protocol to identify parties that abort so that if an abort occurs the faulty party can be excluded and the others start again without them.

We note that all these protocols assume that classical MPC is available as a resource,usually for the same number of dishonest players and the same abort conditions as the constructed MPQC. Therefore, all these results require the same setup assumptions as the corresponding classical MPC. For example, for $t < n/3$ and guaranteed output, one can perform classical MPC assuming only pairwise secure channels between the players (Cramer, Damgård, and Nielsen, 2015). For $t < n/2$ and guaranteed output, one additionally needs to broadcast for information-theoretic security, but pairwise authentic channels are sufficient for computational security (Cramer, Damgård, and Nielsen, 2015). If we drop fairness, then in the case of computational security one gets unfair security for any $t < n$ if one assumes oblivious transfer (Goldreich, Micali, and Wigderson, 1987) or a common reference string (Canetti *et al.*, 2002), and information-theoretic security if one assumes common shared randomness (Ishai, Ostrovsky, and Zikas, 2014).

6. One-time programs

A special class of multiparty functionalities that have been studied in more detail are nonreactive, sender-oblivious functions; i.e., one player is labeled "sender" and another "receiver," and only the receiver obtains the output of the function. This special structure allows for noninteractive protocols to construct a resource that computes such a function: communication goes only from the sender to the receiver. The receiver can use the information obtained to evaluate the function on one input. But, by definition of the ideal resource, he may not repeat this on a second input. The corresponding resources are sometimes called one-time programs. Goyal *et al.* (2010) gave a construction for one-time programs that starts, however, with a resource that is similar to oblivious transfer.[55] This has been called one-time memory or a hardware token (Goyal *et al.*, 2010) since it could be implemented given hardware assumptions, such as a one-time memory that contains the two strings s_0 and s_1, but self-destructs after producing an output.

These results were generalized to the quantum setting by Broadbent, Gutoski, and Stebila (2013), who showed that one can construct quantum one-time programs given access to the same one-time memory resources as for classical one-time programs. More precisely, Broadbent, Gutoski, and Stebila (2013) showed that any completely positive, trace-preserving map $\Phi: \mathcal{H}_A \otimes \mathcal{H}_B \to \mathcal{H}_C$ can be evaluated with a noninteractive protocol by two distrustful parties holding the inputs of registers A and B, respectively, provided that only one player is expected to receive the output in register C.

E. Relativistic cryptography

Thus far we have predominantly discussed protocols whose security is based on the laws of quantum theory. One may, however, further exploit physical laws, such as those of special relativity. These imply an upper bound on the velocity by which information can spread: the velocity of light. The combination of quantum information theory and relativity, apart from its relevance to fundamental questions (Peres and Terno, 2004), opens the possibility of achieving certain cryptographic tasks that are provably impossible based on quantum theory alone.

An example of this is relativistic bit commitment (Kent, 1999), which we mentioned in Sec. VIII.D.1. Another one is coin flipping. Here the two players Alice and Bob each have a trusted agent at two locations L_1 and L_2. At L_1, agent A_1 is instructed to provide agent B_1 with a random bit a, and at L_2 agent B_2 provides A_2 with a random bit b. The agents then inform Alice and Bob about these values, who can then output $a \oplus b$ as the result of the coin flip. The distance between the locations L_1 and L_2 must be chosen large enough to ensure that, if Bob is cheating, he cannot wait until he learns a and then choose b depending on that value. Alice cannot cheat for the same reason. The output $a \oplus b$ is hence uniformly random, provided that at least one of the players chooses their bit uniformly at random.

This protocol does not, however, construct a coin flip resource, because the parallel execution of two instances of the protocol does not behave identically to two coin flip resources in parallel. Suppose that Alice and Bob are running the protocol as well as Bob and Charlie, who send their agents to the same locations L_1 and L_2. At L_1, A_1 then gives her random bit a to B_1, who in turn gives a copy to C_1. And at L_2, C_2 gives his bit c to B_2, who in turn gives a copy to A_2. Alice and Charlie then end up with exactly the same coin flip $a \oplus c$. But, if we were to run two coin flip resources in parallel, we would obtain two independent bit flips.

Running the same kind of attack on the relativistic bit commitment protocols (Kent, 1999, 2012; Kaniewski *et al.*, 2013), Bob may forward Alice's commitment to Charlie and convince Charlie that he is committed to a known bit, whereas in reality he does not know the commitment, and thus does not satisfy the requirement of the bit commitment resource. The same principle was used by Brassard *et al.* (1998) to prove that some noncomposable definitions of bit commitment appearing in the classical literature cannot be used to force a quantum player to commit to a measurement outcome. This technique was then used in Appendix A of Kaniewski (2015) to prove that composing the previously mentioned relativistic bit commitment protocols with the oblivious transfer protocol of Unruh (2010) is insecure. More precisely, in the attack of Brassard *et al.* (1998) and Kaniewski (2015), the committer does not measure her state as required by the protocol, but instead runs the protocol in superposition and measures only the strings that she needs to send to the receiver as part of the commitment protocol. If she is asked to open the commitment, she can still perform the required measurement and open correctly. But if she is not asked to open, she can "undo" this measurement and recover the original state.

Relativistic bit commitment (and coin flipping) was analyzed more systematically by Vilasini, Portmann, and Rio (2019) and Prokop (2020) using the abstract cryptography framework (Maurer and Renner, 2011). More precisely, they instantiated the systems model from Maurer and Renner (2011) with the causal boxes framework (Portmann *et al.*, 2017), which can model information with positions in space-time. The resulting framework was used to prove both impossibility and possibility results for relativistic cryptography.

Vilasini, Portmann, and Rio (2019) showed that it is impossible to construct a biased coin flip resource between two distrustful players without assuming any resources to help them, even in a relativistic setting. Since such a biased coin flip can be constructed from bit commitment (Blum, 1983; Demay and Maurer, 2013), this immediately implies that it is also impossible to construct a bit commitment resource in a relativistic setting. The impossibility results also hold against adversaries that are computationally bounded or have bounded storage (Vilasini, Portmann, and Rio, 2019). Prokop (2020) analyzed what extra resources one can assume to have in the real world to get positive results. The techniques of Vilasini, Portmann, and Rio (2019) were extended by Laneve and Rio (2021) to prove that oblivious transfer is also impossible without setup assumptions other than communication channels, even if the adversary is computationally bounded and has bounded or noisy quantum storage.

[55] Note that oblivious transfer allows the player preparing the two strings s_0 and s_1 to learn whether the other player has queried s_c. With only one-way communication of one-time programs, the resource used cannot allow this, but otherwise it is identical to an oblivious transfer resource.

Another task for which special relativity is taken into account is position verification: a prover wishes to convince a verifier that she is in a specific location (Chandran *et al.*, 2009). Protocols based on relativity have been designed for other tasks as well, such as position verification and authentication (Buhrman *et al.*, 2014; Unruh, 2014). Buhrman *et al.* (2014) showed that such a task is impossible in the presence of multiple colluding quantum adversaries that share entanglement. They also considered a model in which holding shared entanglement is not allowed and proposed a protocol for position verification in this model. Similarly, Unruh (2014) proposed a protocol for position verification in the random oracle model, but with no restriction on entanglement or memory. But neither of these results provides a composable security proof, so it remains open to prove exactly what these protocols achieve; see also Sec. IX.C.

F. Secure quantum message transmission with computational security

Most of the quantum cryptography literature is dedicated to information-theoretic security since the main motivation of this field of research is to reduce cryptographic security to physical principles. This means that regardless of the computational ability of the adversary, the scheme may not be broken, as it does not leak any information about the key or message. It is nevertheless sensible to consider computational security, as certain cryptographic tasks may only be possible under such restricted security guarantees; see Alagic *et al.* (2016) and the references therein. In such a paradigm, a scheme may not be broken by an adversary that is computationally bounded, but with unlimited computational power it may be possible to obtain secret keys or read private messages.

Composable frameworks such as those given by Pfitzmann and Waidner (2000, 2001) and Canetti (2001) and the quantum version by Unruh (2010) all define both computational and information-theoretic security. However, they define security only asymptotically. In other words, a protocol is parametrized by some security parameter n (typically, this might correspond to the number of signals exchanged between the parties or length of a tag appended to a message) and security is proven in the limit as $n \to \infty$. Abstract cryptography (Maurer and Renner, 2011), on the other hand, considers finite security; i.e., a security statement is made for every n (the limit is ignored and may not even be well defined).

Since any implementation is necessarily finite (the players fix a value for the security parameter n that they consider to be sufficient and implement the corresponding protocol) an asymptotic statement is arguably of limited interest in practice. For this reason a paradigm known as concrete security was proposed (Bellare *et al.*, 1997) in which parameters and reductions are given explicitly instead of being hidden in O notation and polynomial-time statements. This allows a user to recover exact bounds for every n instead of being provided with only the limit values.

Concrete security, however, is still intrinsically asymptotic since adversaries are required to be polynomial time in n; protocols, reductions, and simulators need to be efficient in n; errors have to be negligible in n; and such concepts are all defined asymptotically. In finite security, one analyzes the security of a protocol for individual values $n = n_0$. Hence, concepts such as polynomial time, efficiency, or negligibility are not necessarily well defined in a finite analysis and cannot be part of a security definition.

In Sec. VIII.F.1 we explain how to define finite computational security. This follows the paradigm of AC (Maurer and Renner, 2011; Maurer, 2012; Maurer and Renner, 2016) and was used by Maurer, Rüedlinger, and Tackmann (2012), Coretti, Maurer, and Tackmann (2013), and Banfi *et al.* (2019). In Sec. VIII.F.2 we review the results of Banfi *et al.* (2019) on computational security of quantum message transmission (QMT). And in Sec. VIII.F.3 we discuss some asymptotic game-based security definitions for QMT that have been proposed in the literature (Alagic, Gagliardoni, and Majenz, 2018).

1. Defining composable and finite computational security

To adapt the framework described in Sec. II to capture computational security, one needs to change the pseudometric used to distinguish systems. We first recall Eqs. (4) and (6) from Sec. II.F, namely, that the distinguishing advantage for a distinguisher $\mathfrak{D}$ is defined as

$$d^{\mathfrak{D}}(\mathcal{R},\mathcal{S}) \coloneqq |\Pr[\mathfrak{D}(\mathcal{R})=0] - \Pr[\mathfrak{D}(\mathcal{S})=0]|, \quad (38)$$

and the distinguishing advantage for a class of distinguishers $\mathbb{D}$ is given by

$$d^{\mathbb{D}}(\mathcal{R},\mathcal{S}) \coloneqq \sup_{\mathfrak{D}\in\mathbb{D}} d^{\mathfrak{D}}(\mathcal{R},\mathcal{S}). \quad (39)$$

Thus far in this work we have taken $\mathbb{D}$ to be the set of all distinguishers. If a protocol is only computationally secure, Eq. (39) could be large since some ("inefficient") distinguisher might be able to distinguish between the real and ideal systems. Thus, instead of bounding the supremum over all distinguishers as in Eq. (39), we bound Eq. (38) for all $\mathfrak{D}$; i.e., one needs to find some function $f: \mathbb{D} \to \mathbb{R}$ such that,[56] for all $\mathfrak{D}$,

$$d^{\mathfrak{D}}(\mathcal{R},\mathcal{S}) \le f(\mathfrak{D}). \quad (40)$$

Typically, such a bound is given by a reduction; i.e., one proves that if a distinguisher $\mathfrak{D}$ can distinguish the real from the ideal system, then $\mathfrak{D}$ may be used to solve a problem that is believed to be hard. In Eq. (40), $f(\mathfrak{D})$ may then correspond to the probability that this problem may be solved using $\mathfrak{D}$.[57]

[56]In asymptotic security one may still use Eq. (39) instead of Eq. (40) but may replace $\mathbb{D}$ with the set of all efficient distinguishers. This is not well defined in the finite setting, since efficiency is defined only asymptotically.

[57]A detailed explanation of this paradigm for modeling computational security using a reduction to computationally hard problems was provided by Rogaway (2006) within a classical asymptotic model.

Note that information-theoretic security corresponds to the special case in which one can show that $f(\mathfrak{D})$ is small for all $\mathfrak{D}$; i.e., a security proof with error ε means that Eq. (40) is bounded by $f(\mathfrak{D}) = \varepsilon$ for all $\mathfrak{D}$.

For a longer exposition on finite computational security, see Banfi *et al.* (2019).

2. Secure quantum message transmission

To construct a secure quantum channel with information-theoretic security, the secret key shared by the honest players needs to be longer than twice the length of quantum message sent; see Sec. VIII.A. Although QKD or key recycling (see Sec. VIII.B) may be used to obtain more keys, this only works when the noise on the channel is sufficiently low or the adversary decides not to tamper with the messages. Should the noise be too high, the used key is irremediably lost and the honest players may run out of keys and no longer be able to communicate securely.

With computational security, it is believed[58] that the key can be much shorter than the message, essentially allowing the same key to be used over and over, even when conditions do not allow for recycling. For example, it is believed that one can construct (quantum resistant) pseudorandom function (PRF) families (Zhandry, 2012). In other words, there are families of functions $\{f_k: \{0,1\}^m \to \{0,1\}^n\}_{k\in\mathcal{K}}$ such that if k is chosen uniformly at random then the output of f_k cannot be distinguished by a computationally bounded player from a random oracle (RO) that outputs a uniform random string for every new input. If Alice and Bob share a secret key k that they use to pick a function f_k, then they may encrypt their first message using $f_k(1)$ as key, encrypt their second message using $f_k(2)$ as key, etc. To a computationally bounded adversary not knowing k, the encryption keys $f_k(i)$ would look random, and hence, by composition, a scheme requiring a uniformly random secret key would be secure when used with these pseudorandom keys.

Exactly this was done by Banfi *et al.* (2019), who composed a PRF family with the information-theoretic QMT protocol from Sec. VIII.A.[59] They proved that if the error of the PRF is bounded by

$$d^{\mathfrak{D}}(\mathsf{PRF}, \mathsf{RO}) \leq \varepsilon^{\mathsf{PRF}}(\mathfrak{D}),$$

and if the QMT protocol has error $\varepsilon^{\mathsf{QMT}}$ and is used to send at most ℓ messages, then the composed protocol essentially constructs ℓ copies of the secure channel from Fig. 31. The error of this construction is bounded by

$$\varepsilon(\mathfrak{D}) = \ell\varepsilon^{\mathsf{QMT}} + \varepsilon^{\mathsf{PRF}}(\mathfrak{D}'),$$

where $\mathfrak{D}'$ is the same distinguisher as $\mathfrak{D}$ with the additional quality that it can perform ℓ extra encryption and decryption operations.

3. Relation to other security definitions

Computational security for secure message transmission is often characterized by asymptotic game-based definitions: e.g., an adversary chooses two plain texts, receives a ciphertext for one of the two, and has to guess which of the two plain texts it corresponds to. To model situations in which the same keys can be used to encrypt and decrypt other messages that may be accessible to the adversary, she is also given oracle access to either encryption or decryption functions at various points of the game (Bellare *et al.*, 1998; Katz and Yung, 2006). These definitions were adapted to the quantum case by Broadbent and Jeffery (2015), Alagic *et al.* (2016), and Alagic, Gagliardoni, and Majenz (2018).

Before such definitions may be safely used in practice, it is essential to understand which security guarantees they provide, i.e., which resources they assume and which resources they construct. The accessible information security definition for QKD that was discussed in Sec. III.C.1 (see also Sec. VI.F) turned out to implicitly assume that the adversary has no quantum memory. In the case of these game-based definitions, a series of results showed that some of them have the opposite flaw: they construct a resource that is unnecessarily strong and exclude certain protocols that should be considered secure (Canetti, Krawczyk, and Nielsen, 2003; Coretti, Maurer, and Tackmann, 2013; Banfi *et al.*, 2019).

The strongest of these definitions, called quantum authenticated encryption (QAE) by Alagic, Gagliardoni, and Majenz (2018), is the most similar to the construction of secure channels used in Sec. VIII.F.2. Banfi *et al.* (2019) showed that QAE essentially corresponds to constructing a secure channel, but with a fixed simulator, whereas a security definition within the abstract cryptography framework requires only the existence of a simulator. A protocol for which the simulator hard coded in QAE is a good simulator will be deemed secure. However, for a protocol that requires a different simulator to prove its security, the QAE definition will merely declare it insecure, even though it constructs a secure channel.

IX. OPEN PROBLEMS

Only a relatively small part of the protocols appearing in the quantum cryptography literature has been analyzed and proved secure within a composable framework. To understand the security guarantees they actually provide, and in which contexts they can be used safely, such an analysis would, however, be crucial, and thus represents a major task for quantum cryptographers complete in the future. Here we illustrate this task, focusing on a few areas that we consider interesting. The first is the problem of reusing devices in device-independent cryptography (Sec. IX.A). The second is modeling quantum cryptography with nonasymptotic computational assumptions (Sec. IX.B). And the third consists of studying setup assumptions that are needed to achieve a broader range of constructions (Sec. IX.C).

[58]Computational security always relies on the belief that some problem is hard to solve. A cryptographic security proof then consists of showing that, if an adversary can break the scheme, then this adversary can also solve the hard problem.

[59]A variant of this protocol that allows the adversary to jumble the order of the messages was first proposed by Alagic, Gagliardoni, and Majenz (2018), but security was proven only by using asymptotic game-based definitions; see Sec. VIII.F.3.

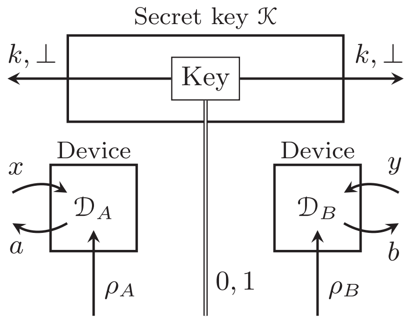

FIG. 34. An ideal world in which a secret key is produced by $\mathcal{K}$ and new devices $\mathcal{D}_A$ and $\mathcal{D}_B$ independent of $\mathcal{K}$ are accessible to the players.

A. Reusing devices in device-independent cryptography

In Sec. VI.D we modeled DI QKD. There the untrusted devices correspond to resources that are available to the honest players. If Alice and Bob want to run another DI-QKD protocol to generate more keys once the first run is over, they will again need all the same resources; i.e., they will need such devices once more. If Alice and Bob have access to new, fresh devices, they can run the protocol a second time with them. However, it does not follow from that analysis that the same devices can be used again. In fact, it was shown by Barrett, Colbeck, and Kent (2013) that in general these devices cannot be reused a second time. The internal memory of a device used for key (or randomness) generation may contain information about the secret key (or random number) generated in the first round, and the device may thus leak this information when it is reused in a second round. A secret bit may be leaked in a subtle manner. For example, if the bit equals 0 the device may perform the expected operations during the second round, and if the bit equals 1 it may force an abort.

Reusing devices in DI cryptography is similar to reusing keys. In general it cannot be done. However, in the case of keys, if one can prove that the key is close to uniform and independent of the adversary's information, then it can be recycled; this was covered in Sec. VIII.B. The same approach could be used to recycle devices: instead of the ideal world consisting simply of a key resource, it should also provide access to devices that are independent of this key, as depicted in Fig. 34.

No DI-QKD protocol has ever been shown to construct the ideal system from Fig. 34, and it may well be impossible to do so. But even if this is the case, it does not exclude the possibility that one can construct an ideal system that is stronger than the shared secret key considered in Sec. VI.D, such as one in which the devices have some partial independence from the key or are fully independent in certain contexts.[60]

We note that weaker models such as MDI QKD (see Sec. VI.E) do not suffer from the same problem of device reuse as DI QKD. The reason is that in MDI QKD one does not need to make any assumptions about the measurement devices (the adversary does the measurements for the honest players), whereas in DI QKD one has to assume that no unauthorized information leaves the devices.

B. Computational security

Computational security is a fairly unexplored area of quantum cryptography. The main motivation for studying this is to achieve results that are not possible with information-theoretic security. For example, in Sec. VIII.C we mentioned a computationally secure protocol for delegated quantum computation with a classical client (Gheorghiu and Vidick, 2019), which is not believed to be possible with information-theoretic security (Aaronson *et al.*, 2019). The computationally secure message transmission from Sec. VIII.F allows keys to be reused without the extra communication required by QKD (Sec. III) or key recycling (Sec. VIII.B), and thus without the possibility of an adversary interrupting this communication and preventing the key from being reused. And the work of Unruh (2013) discussed in Sec. VIII.D.4 removed the need for a shared secret key in QKD by using signature cards instead.

One may essentially analyze any area of cryptography with computational security to study how assumptions needed for information-theoretic security may be weakened in the computational setting. There is, however, no single way to model computational assumptions, and important open questions in the field are to identify the best ways of doing this. Most frequently, one proves a reduction; i.e., if some distinguisher can guess whether it is interacting with the real or ideal system, then this distinguisher can be used to solve some problem that is believed to be hard. In Sec. VIII.F we reviewed the finite reductions of Banfi *et al.* (2019), in which the probability of a distinguisher $\mathfrak{D}$ distinguishing the real and ideal worlds is bounded by the probability of this distinguisher being successfully used (as part of a new distinguisher $\mathfrak{D}'$) to distinguish a pseudorandom function from a uniform random function; see also Rogaway (2006) for a discussion of reductions.

Another way to define computational security would be to define an ideal resource that falls under the control of the adversary if she can solve some problem believed to be hard (such as finding a collision for a hash function). This is essentially the "identical-until-bad" concept of Bellare and Rogaway (2006) but adapted to composable security instead of game-based security. To the best of our knowledge, this paradigm remains completely unexplored in quantum cryptography.

Other works such as that of Chen *et al.* (2017) bound adversaries by circuit sizes. It is not clear how to model that in a finite, composable framework and is important open work.

C. Other setup assumptions

When a security definition is considered "not composable," it often has a setup assumption hard coded in it that is not present in the obvious composable definition, and is therefore strictly weaker. By modeling this assumption in a composable framework, one can get another, equivalent composable definition. We illustrated this in Sec. VI.F by explaining

[60]Context restricted composability is a promising research path for protocols that do not construct the desired ideal resource. Its investigation was initiated by Jost and Maurer (2018) and is beyond the scope of this review.

how a definition for QKD based on the accessible information, which is normally not composable, can be turned into a composable one within a model where an adversary has no long-term quantum memory.

Similar techniques were used by Unruh (2011) to obtain commitments in the bounded storage model. While it follows from Vilasini, Portmann, and Rio (2019) that coin flipping and bit commitment are impossible in a bounded storage model without further assumptions, Unruh (2011) avoided these by putting a bound on the number of times a protocol can be run in parallel, and designing protocols that are secure for this limited number of compositions.[61] Likewise, Prokop (2020) made extra setup assumptions in the relativistic model to avoid the impossibility results of Vilasini, Portmann, and Rio (2019).

There are numerous works where security is proved based on the assumption that adversaries are restricted. For example, adversaries cannot share entanglement in the work of Buhrman *et al.* (2014), the adversaries' memory size is bounded in the work of Damgård *et al.* (2007, 2008), the adversaries' memory is noisy in the work of Wehner, Schaffner, and Terhal (2008), Schaffner, Terhal, and Wehner (2009), and König, Wehner, and Wullschleger (2012), and adversaries can perform only local operations on single qubits and communicate only classically in the work of Liu (2014, 2015). It remains open how to model these assumptions to get composable security statements and prove in what setting such protocols are secure. Similarly, to capture position-based cryptography Unruh (2014) used a model of circuits with positions in space-time. Here too it is not clear how to fit these results into a composable framework and identify the resource that is constructed by these protocols.

ACKNOWLEDGMENTS

We thank Claude Crépeau, Marco Lucamarini, Mark Wilde, Ramona Wolf, and the anonymous referees for their helpful comments and suggestions. Most of the work of C. P. was completed at ETH Zurich. This work has been funded by the European Research Council (ERC) via Grant No. 258932, the Swiss National Science Foundation via the National Centre of Competence in Research "QSIT," the QuantERA project "eDICT," the U.S. Air Force Office of Scientific Research (AFOSR) via Grants No. FA9550-16-1-0245 and No. FA9550-19-1-0202, the Quantum Center of ETH Zurich, and the Zurich Information Security and Privacy Center.

APPENDIX A: TRACE DISTANCE

Many of the statements in this review use the well-known fact that the distinguishing advantage between two systems that output states ρ and σ is equivalent to the trace distance between these states. In this appendix we gather lemmas and theorems that prove this fact and help interpret the meaning of the trace distance.

[61]This effectively restricts what the distinguisher environment may do to distinguish the real and ideal systems since the bound on the number of executions of a protocol applies to the distinguisher as well.

In Appendix A.1 we first define the trace distance (as well as its classical counterpart, the total variation distance) and provide some basic lemmas that can also be found in textbooks such as that of Nielsen and Chuang (2010). In Appendix A.2 we then show the connection between trace distance and distinguishing advantage, which was originally proven by Helstrom (1976). In Appendix A.3 we prove that we can alternatively think of the trace distance between a real and an ideal system as a bound on the probability that a failure occurs in the real system, as suggested by Renner (2005). Finally, in Appendix A.4 we bound two typical information theory notions of secrecy (the conditional entropy of a key given the eavesdropper's information and her probability of correctly guessing the key) in terms of the trace distance. Although such measures of information are generally ill suited for defining cryptographic security, they can help interpret the notion of a key being ε close to uniform. We refer the interested reader to Trushechkin (2020) for further interpretations of the trace distance.

1. Metric definitions

In the case of a classical system, statistical security is defined by the total variation (or statistical) distance between the probability distributions describing the real and ideal settings, which is defined as follows.[62]

Definition 4 (total variation distance).—The total variation distance between two probability distributions P_Z and $P_{\tilde{Z}}$ over an alphabet $\mathcal{Z}$ is defined as

$$D(P_Z, P_{\tilde{Z}}) \coloneqq \frac{1}{2}\sum_{z\in\mathcal{Z}}|P_Z(z) - P_{\tilde{Z}}(z)|.$$

Using the fact that $|a-b| = a+b-2\ \min(a,b)$, one can also write the total variation distance as

$$D(P_Z, P_{\tilde{Z}}) = 1 - \sum_{z\in\mathcal{Z}}\min[P_Z(z), P_{\tilde{Z}}(z)]. \tag{A1}$$

In the case of quantum states instead of classical random variables, the total variation distance generalizes to the trace distance. More precisely, the trace distance between two density operators that are diagonal in the same orthonormal basis is equal to the total variation distance between the probability distributions defined by their respective eigenvalues.

Definition 5 (trace distance).—The trace distance between two quantum states ρ and σ is defined as

$$D(\rho,\sigma) \coloneqq \tfrac{1}{2}\mathrm{tr}|\rho-\sigma|.$$

We now introduce some technical lemmas involving the trace distance that help us derive the theorems to follow. Proofs of these lemmas were given by Nielsen and Chuang (2010).

Lemma 6.—For any two states ρ and σ and any operator $0 \le M \le I$,

[62]We employ the same notation $D(\cdot,\cdot)$ for both the total variation and the trace distance since the former is a special case of the latter.

$$D(\rho,\sigma) \geq \mathrm{tr}[M(\rho-\sigma)]. \tag{A2}$$

Furthermore, this inequality is tight for some values of M.

The trace distance can thus be alternatively written as

$$D(\rho,\sigma) = \max_M \mathrm{tr}[M(\rho-\sigma)]. \tag{A3}$$

Let $\{\Gamma_x\}_x$ be a POVM) (a set of operators $0 \leq \Gamma_x \leq I$ such that $\sum_x \Gamma_x = I$), and let P_X denote the outcome of measuring a quantum state ρ with $\{\Gamma_x\}_x$, i.e., $P_X(x) = \mathrm{tr}(\#2)\Gamma_x\rho$. Lemma 7 says that the trace distance between two states ρ and σ is equal to the total variation between the outcomes (P_X and Q_X) of an optimal measurement on the two states.

Lemma 7.—For any two states ρ and σ,

$$D(\rho,\sigma) = \max_{\{\Gamma_x\}_x} D(P_X, Q_X), \tag{A4}$$

where P_X and Q_X are the probability distributions resulting from measuring ρ and σ, respectively, with a POVM $\{\Gamma_x\}_x$ and the maximization is over all POVMs. Furthermore, if the two states ρ_{ZB} and σ_{ZB} have a classical subsystem Z, then the measurement satisfying Eq. (A4) leaves the classical subsystem unchanged; i.e., the maximum is reached for a POVM with elements

$$\Gamma_x = \sum_z |z\rangle\langle z| \otimes M_x^z, \tag{A5}$$

where $\{|z\rangle\}_z$ is the classical orthonormal basis of Z.

2. Distinguishing advantage

Helstrom (1976) proved that the advantage a distinguisher has in guessing whether it was provided with one of two states with equal priors ρ or σ is given by the trace distance[63] between the two $D(\rho,\sigma)$. We first sketch the classical case, then prove the quantum version.

Let a distinguisher be given a value sampled according to probability distributions P_Z or $P_{\tilde{Z}}$, where P_Z and $P_{\tilde{Z}}$ are each chosen with probability $1/2$. Suppose that the value received by the distinguisher is $z \in \mathcal{Z}$. If $P_Z(z) > P_{\tilde{Z}}(z)$, its best guess is that the value was sampled according to P_Z. Otherwise, it should guess that it was $P_{\tilde{Z}}$. Let $\mathcal{Z}' \coloneqq \{z \in \mathcal{Z} : P_Z(z) > P_{\tilde{Z}}(z)\}$ and $\mathcal{Z}'' \coloneqq \{z \in \mathcal{Z} : P_Z(z) \leq P_{\tilde{Z}}(z)\}$. There are a total of $2|\mathcal{Z}|$ possible events: the sample is chosen according to P_Z or $P_{\tilde{Z}}$ and takes the value $z \in \mathcal{Z}$. These events have probabilities $P_Z(z)/2$ and $P_{\tilde{Z}}(z)/2$. Conditioned on P_Z being chosen and z being the sampled value, the distinguisher has probability 1 of guessing correctly with the previously outlined strategy if $z \in \mathcal{Z}'$, and 0 otherwise. Likewise, if $P_{\tilde{Z}}$ is selected, it has probability 1 of guessing correctly if $z \in \mathcal{Z}''$, and 0 otherwise. The probability of correctly guessing whether it was given a value sampled according to P_Z or $P_{\tilde{Z}}$, which we denote as $p_{\text{distinguish}}(P_Z, P_{\tilde{Z}})$, is obtained by summing over all possible events weighted by their probabilities. Hence,

$$\begin{aligned}
p_{\text{distinguish}}(P_Z, P_{\tilde{Z}}) &= \sum_{z\in\mathcal{Z}'} \frac{P_Z(z)}{2} + \sum_{z\in\mathcal{Z}''} \frac{P_{\tilde{Z}}(z)}{2} \\
&= \frac{1}{2}\left(1 - \sum_{z\in\mathcal{Z}''} P_Z(z)\right) + \frac{1}{2}\left(1 - \sum_{z\in\mathcal{Z}'} P_{\tilde{Z}}(z)\right) \\
&= 1 - \frac{1}{2}\sum_{z\in\mathcal{Z}} \min[P_Z(z), P_{\tilde{Z}}(z)] \\
&= \tfrac{1}{2} + \tfrac{1}{2} D(P_Z, P_{\tilde{Z}}),
\end{aligned}$$

where in the last equality we used the alternative formulation of the total variation distance from Eq. (A1).

We now generalize the previous argument to quantum states with equal priors, which is a special case given by Helstrom (1976).

Theorem 8.—For any states ρ and σ, we have

$$p_{\text{distinguish}}(\rho,\sigma) = \tfrac{1}{2} + \tfrac{1}{2} D(\rho,\sigma).$$

Proof.—If a distinguisher is given one of two states ρ or σ, each with probability $1/2$, its probability of guessing which one it holds is given by a maximization of all possible measurements it may do.It chooses a POVM $\{\Gamma_0, \Gamma_1\}$, where Γ_0 and Γ_1 are positive operators with $\Gamma_0 + \Gamma_1 = I$, and measures the state it holds. If it gets the outcome 0, it guesses that it holds ρ and, if it gets the outcome 1, it guesses that it holds σ. The probability of guessing correctly is given by

$$\begin{aligned}
p_{\text{distinguish}}(\rho,\sigma) &= \max_{\Gamma_0,\Gamma_1} [\tfrac{1}{2}\mathrm{tr}(\Gamma_0\rho) + \tfrac{1}{2}\mathrm{tr}(\Gamma_1\sigma)] \\
&= \tfrac{1}{2}\max_{\Gamma_0} \{\mathrm{tr}(\Gamma_0\rho) + \mathrm{tr}[(I-\Gamma_0)\sigma]\} \\
&= \tfrac{1}{2} + \tfrac{1}{2}\max_{\Gamma_0} \mathrm{tr}[\Gamma_0(\rho-\sigma)].
\end{aligned} \tag{A6}$$

The proof concludes by plugging Eq. (A3) in Eq. (A6). ■

3. Probability of a failure

The trace distance is used as the security definition of QKD because the relevant measure for cryptographic security is the distinguishing advantage (as discussed in Sec. II), and as proven in Theorem 8 the distinguishing advantage between two quantum states corresponds to their trace distance. This operational interpretation of the trace distance involves two worlds, an ideal one and a real one, and the distance measure is the renormalized difference between the probabilities of the distinguisher correctly guessing which world it is connected to.

In this section we describe a different interpretation of the total variation and trace distances. Instead of having two different worlds, we consider one world in which the outcomes of interacting with the real and ideal systems coexist. And, instead of these distance measures being a difference between probability distributions, they become the probability that any classical value occurring in one of the systems does not simultaneously occur in the other. We call such an event a

[63] Actually, Helstrom (1976) solved a more general problem, in which the states ρ and σ are picked with *a priori* probabilities p and $1-p$, respectively, instead of $1/2$ as in the definition of the distinguishing advantage.

failure (since one system is ideal, if the other behaves differently, it must have failed) and the trace distance becomes the probability of a failure occurring.

Given two random variables Z and $\tilde{Z}$ with probability distributions P_Z and $P_{\tilde{Z}}$, any distribution $P_{Z\tilde{Z}}$ with marginals given by P_Z and $P_{\tilde{Z}}$ is called a coupling of P_Z and $P_{\tilde{Z}}$. The interpretation of the trace distance treated in this section uses one specific coupling, known as a maximal coupling in probability theory (Thorisson, 2000).

Theorem 9 (maximal coupling).—Let P_Z and $P_{\tilde{Z}}$ be two probability distributions over the same alphabet $\mathcal{Z}$. There is then a probability distribution $P_{Z\tilde{Z}}$ on $\mathcal{Z}\times\mathcal{Z}$ such that

$$\Pr[Z=\tilde{Z}] \coloneqq \sum_z P_{Z\tilde{Z}}(z,z) \geq 1 - D(P_Z, P_{\tilde{Z}}), \tag{A7}$$

and such that P_Z and $P_{\tilde{Z}}$ are the marginals of $P_{Z\tilde{Z}}$, i.e.,

$$P_Z(z) = \sum_{\tilde{z}} P_{Z\tilde{Z}}(z,\tilde{z}) \quad (\forall\ z\in\mathcal{Z}), \tag{A8}$$

$$P_{\tilde{Z}}(\tilde{z}) = \sum_{z} P_{Z\tilde{Z}}(z,\tilde{z}) \quad (\forall\ \tilde{z}\in\mathcal{Z}). \tag{A9}$$

It turns out that the inequality in Eq. (A7) is tight, i.e., one can also show that, for any distribution $P_{Z\tilde{Z}}$, $\Pr[Z=\tilde{Z}] \leq 1 - D(P_Z, P_{\tilde{Z}})$. We will not, however, use this fact here.

Consider now a real system that outputs values given by Z and an ideal system that outputs values according to $\tilde{Z}$. Theorem 9 tells us that there is a coupling of these distributions such that the probability of the real system producing a different value from the ideal system is bounded by the total variation distance between P_Z and $P_{\tilde{Z}}$. Thus, the real system behaves ideally except with probability $D(P_Z, P_{\tilde{Z}})$.

We first prove this theorem, then in Corollary 10 apply it to quantum systems.

Proof of Theorem 9.—Let $Q_{Z\tilde{Z}}$ be the real function on $\mathcal{Z}\times\mathcal{Z}$ defined by

$$Q_{Z\tilde{Z}}(z,\tilde{z}) = \begin{cases} \min[P_Z(z), P_{\tilde{Z}}(\tilde{z})] & \text{if } z=\tilde{z}, \\ 0 & \text{otherwise} \end{cases}$$

(for all $z,\tilde{z}\in\mathcal{Z}$). Furthermore, let R_Z and $R_{\tilde{Z}}$ be the real functions on $\mathcal{Z}$ defined by

$$R_Z(z) = P_Z(z) - Q_{Z\tilde{Z}}(z,z),$$
$$R_{\tilde{Z}}(\tilde{z}) = P_{\tilde{Z}}(\tilde{z}) - Q_{Z\tilde{Z}}(\tilde{z},\tilde{z}).$$

We then define $P_{Z\tilde{Z}}$ by

$$P_{Z\tilde{Z}}(z,\tilde{z}) = Q_{Z\tilde{Z}}(z,\tilde{z}) + \frac{1}{D(P_Z,P_{\tilde{Z}})} R_Z(z) R_{\tilde{Z}}(\tilde{z}).$$

We now show that $P_{Z\tilde{Z}}$ satisfies the conditions of the theorem. For this, we note that for any $z\in\mathcal{Z}$

$$R_Z(z) = P_Z(z) - \min[P_Z(z), P_{\tilde{Z}}(z)] \geq 0.$$

That is, R_Z and, likewise, $R_{\tilde{Z}}$ are non-negative. Since $Q_{Z\tilde{Z}}$ is by definition also non-negative, we find that $P_{Z\tilde{Z}}$ is non-negative too. From Eq. (A8) or (A9), which we later prove, it follows that $P_{Z\tilde{Z}}$ is also normalized. Hence, $P_{Z\tilde{Z}}$ is a valid probability distribution.

To show Eq. (A7), we again use the non-negativity of R_Z and $R_{\tilde{Z}}$, which implies

$$\begin{aligned}\sum_z P_{Z\tilde{Z}}(z,z) &\geq \sum_z Q_{Z\tilde{Z}}(z,z) \\ &= \sum_z \min[P_Z(z), P_{\tilde{Z}}(z)] \\ &= 1 - D(P_Z, P_{\tilde{Z}}),\end{aligned}$$

where in the last equality we used the alternative formulation of the total variation distance from Eq. (A1).

To prove Eq. (A8), we first note that

$$\begin{aligned}\sum_{\tilde{z}} R_{\tilde{Z}}(\tilde{z}) &= \sum_{\tilde{z}} P_{\tilde{Z}}(\tilde{z}) - \sum_{\tilde{z}} Q_{Z\tilde{Z}}(\tilde{z},\tilde{z}) \\ &= 1 - \sum_{\tilde{z}} \min[P_Z(\tilde{z}), P_{\tilde{Z}}(\tilde{z})] \\ &= D(P_Z, P_{\tilde{Z}}).\end{aligned}$$

Using this we find that for any $z\in\mathcal{Z}$

$$\begin{aligned}\sum_{\tilde{z}} P_{Z\tilde{Z}}(z,\tilde{z}) &= \sum_{\tilde{z}} Q_{Z\tilde{Z}}(z,\tilde{z}) + R_Z(z)\frac{1}{D(P_Z,P_{\tilde{Z}})}\sum_{\tilde{z}} R_{\tilde{Z}}(\tilde{z}) \\ &= Q_{Z\tilde{Z}}(z,z) + R_Z(z) = P_Z(z).\end{aligned}$$

By symmetry, this also proves Eq. (A9). ■

In the case of quantum states, Theorem 9 can be used to couple the outcomes of any observable applied to the quantum systems.

Corollary 10.—For any states ρ and σ with trace distance $D(\rho,\sigma) \leq \varepsilon$ and any measurement given by its POVM operators $\{\Gamma_w\}_w$ with outcome probabilities $P_W(w) = \mathrm{tr}(\Gamma_w \rho)$ and $P_{\tilde{W}}(w) = \mathrm{tr}(\Gamma_w \sigma)$, there is a coupling of P_W and $P_{\tilde{W}}$ such that

$$\Pr[W \neq \tilde{W}] \leq D(\rho,\sigma).$$

Proof.—Immediate by combining Lemma 7 and Theorem 9. ■

Corollary 10 tells us that if two systems produce states ρ and σ, then for any observations made on those systems there is a coupling for which the values of each measurement will differ with probability at most $D(\rho,\sigma)$. It is instructive to remember that this operational meaning is not essential to the security notion or part of the framework in any way. It is an intuitive way of understanding the trace distance so as to better choose a suitable value. It allows this distance to be thought of as a maximum failure probability and enables the value for ε to be chosen accordingly.

4. Measures of uncertainty

Noncomposable security models often use measures of uncertainty to quantify how much information an adversary might have about a secret, such as entropy as used by Shannon to prove the security of the one-time pad (Shannon, 1949). These measures are often weaker than what one obtains using a global distinguisher and in general do not provide good security definitions. They are, however, quite intuitive and, in order to further illustrate the quantitative value of the distinguishing advantage, we derive bounds on two of these measures of uncertainty in terms of the trace distance, namely, on the probability of guessing the secret key in Appendix A.4.a and on the von Neumann entropy of the secret key in Appendix A.4.b.

a. Probability of guessing

Let $\rho_{KE} = \sum_{k\in\mathcal{K}} p_k |k\rangle\langle k|_K \otimes \rho_E^k$ be the joint state of a secret key in the K subsystem and Eve's information in the E subsystem. To guess the value of the key, Eve can pick a POVM $\{\Gamma_k\}_{k\in\mathcal{K}}$, measure her system, and output the result of the measurement. Given that the key is k, her probability of having guessed correctly is $\mathrm{tr}(\Gamma_k \rho_E^k)$. The average probability of guessing correctly for this measurement is then given by the sum over all k, weighted by their respective probabilities p_k. And Eve's probability of correctly guessing the key is defined by taking the following maximum over all measurements:

$$p_{\mathrm{guess}}(K|E)_\rho \coloneqq \max_{\{\Gamma_k\}} \sum_{k\in\mathcal{K}} p_k \mathrm{tr}(\Gamma_k \rho_E^k). \qquad \text{(A10)}$$

Lemma 11.—For any bipartite state ρ_{KE} with classical K,

$$p_{\mathrm{guess}}(K|E)_\rho \le \frac{1}{|\mathcal{K}|} + D(\rho_{KE}, \tau_K \otimes \rho_E),$$

where τ_K is the fully mixed state.

Proof.—Note that for $M \coloneqq \sum_k |k\rangle\langle k| \otimes \Gamma_k$, where $\{\Gamma_k\}$ maximizes Eq. (A10), the guessing probability can equivalently be written as

$$p_{\mathrm{guess}}(K|E)_\rho = \mathrm{tr}(M\rho_{KE}).$$

Furthermore,

$$\mathrm{tr}[M(\tau_K \otimes \rho_E)] = \frac{1}{|\mathcal{K}|}.$$

In Lemma 6 we proved that, for any operator $0 \le M \le I$,

$$\mathrm{tr}[M(\rho - \sigma)] \le D(\rho, \sigma).$$

Setting $\rho = \rho_{KE}$ and $\sigma = \tau_K \otimes \rho_E$ in the previous inequality, we finish the proof as follows:

$$\mathrm{tr}(M\rho_{KE}) \le \mathrm{tr}(M)(\tau_K \otimes \rho_E) + D(\rho_{KE}, \tau_K \otimes \rho_E),$$
$$\Rightarrow p_{\mathrm{guess}}(K|E)_\rho \le \frac{1}{|\mathcal{K}|} + D(\rho_{KE}, \tau_K \otimes \rho_E).$$

■

b. Entropy

Let $\rho_{KE} = \sum_{k\in\mathcal{K}} p_k |k\rangle\langle k|_K \otimes \rho_E^k$ be the joint state of a secret key in the K subsystem and Eve's information in the E subsystem. We want to bound the von Neumann entropy of K given $E[S(K|E)_\rho = S(\rho_{KE}) - S(\rho_E)$, where $S(\rho) \coloneqq -\mathrm{tr}(\rho \log \rho)]$ in terms of the trace distance $D(\rho_{KE}, \tau_K \otimes \rho_E)$. We first derive a lower bound on the von Neumann entropy by using the following theorem from Alicki and Fannes (2004).

Theorem 12 [from Alicki and Fannes (2004)].—For any bipartite states ρ_{AB} and σ_{AB} with trace distance $D(\rho, \sigma) = \varepsilon \le 1/4$ and $\dim \mathcal{H}_A = d_A$, we have

$$|S(A|B)_\rho - S(A|B)_\sigma| \le 8\varepsilon \log d_A + 2h(2\varepsilon),$$

where $h(p) = -p \log p - (1-p)\log(1-p)$ is the binary entropy.

Corollary 13.—For any state ρ_{KE} with $D(\rho_{KE}, \tau_K \otimes \rho_E) = \varepsilon \le 1/4$, where τ_K is the fully mixed state, we have

$$S(K|E)_\rho \ge (1 - 8\varepsilon) \log |\mathcal{K}| - 2h(2\varepsilon).$$

Proof.—Immediate by plugging ρ_{KE} and $\tau_K \otimes \rho_E$ into Theorem 12. ■

Given the von Neumann entropy of K conditioned on E, $S(K|E)_\rho$, one can also set an upper bound on the trace distance of ρ_{KE} from $\tau_K \otimes \rho_E$ by relating $S(K|E)_\rho$ to the relative entropy of ρ_{KE} to $\tau_K \otimes \rho_E$.The relative entropy of ρ to σ is defined as $S(\rho\|\sigma) \coloneqq \mathrm{tr}(\rho \log \rho) - \mathrm{tr}(\rho \log \sigma)$.

Lemma 14.—For any quantum state ρ_{KE},

$$D(\rho_{KE}, \tau_K \otimes \rho_E) \le \sqrt{\tfrac{1}{2}[\log |\mathcal{K}| - S(K|E)_\rho]}.$$

Proof.—From the definitions of the relative and von Neumann entropies we have

$$\begin{aligned} S(\rho_{KE}\|\tau_K \otimes \rho_E) &= \log |\mathcal{K}| + S(\rho_{KE}\|\mathrm{id}_K \otimes \rho_E) \\ &= \log |\mathcal{K}| - S(K|E)_\rho, \end{aligned}$$

where id_K is the identity matrix. We then use the following bound on the relative entropy [Theorem 1.15 of Ohya and Petz (1993)] to conclude the proof:

$$S(\rho\|\sigma) \ge 2[D(\rho, \sigma)]^2.$$

■

Corollary 13 and Lemma 14 can be written together as follows in one equation, upper and lower bounding the conditional von Neumann entropy:

$$(1 - 8\varepsilon) \log |\mathcal{K}| - 2h(2\varepsilon) \le S(K|E)_\rho \le \log |\mathcal{K}| - 2\varepsilon^2,$$

where $\varepsilon = D(\rho_{KE}, \tau_K \otimes \rho_E)$.

APPENDIX B: PROOFS FROM SEC. III

In Sec. III we show how to define the security of QKD in a composable framework and relate this to the trace distance security criterion introduced by Renner (2005). This composable treatment of the security of QKD follows the literature (Ben-Or *et al.*, 2005; Müller-Quade and Renner, 2009), and the results presented in Sec. III were given by Ben-Or *et al.* (2005) and Müller-Quade and Renner (2009) as well. The formulation of the statements differs, however, from those works since we use here the abstract cryptography framework of Maurer and Renner (2011). For completeness we provide proofs of the main results from Sec. III here.

Proof of Theorem 2.—Recall that in Sec. III.B.2 we fixed the simulator and showed that to satisfy Eq. (8) it is sufficient for Eq. (11) to hold. Here we break Eq. (11) into security [Eq. (13)] and correctness [Eq. (14)], thus proving the theorem.

We define γ_{ABE} as a state obtained from $\rho_{ABE}^{\top}$ [Eq. (12)] by throwing away the B system and replacing it with a copy of A, i.e.,

$$\gamma_{ABE} = \frac{1}{1-p^{\perp}} \sum_{k_A,k_B\in\mathcal{K}} p_{k_A,k_B} |k_A,k_A\rangle\langle k_A,k_A| \otimes \rho_E^{k_A,k_B}.$$

From the triangle inequality we get

$$D(\rho_{ABE}^{\top}, \tau_{AB}\otimes\rho_E^{\top}) \le D(\rho_{ABE}^{\top}, \gamma_{ABE}) + D(\gamma_{ABE}, \tau_{AB}\otimes\rho_E^{\top}).$$

Since in the states γ_{ABE} and $\tau_{AB}\otimes\rho_E^{\top}$ the B system is a copy of the A system, it does not modify the distance. Furthermore, $\mathrm{tr}_B(\gamma_{ABE}) = \mathrm{tr}_B(\rho_{ABE}^{\top})$. Hence,

$$D(\gamma_{ABE}, \tau_{AB}\otimes\rho_E^{\top}) = D(\gamma_{AE}, \tau_A\otimes\rho_E^{\top}) = D(\rho_{AE}^{\top}, \tau_A\otimes\rho_E^{\top}).$$

For the other term note that

$$\begin{aligned} D(\rho_{ABE}^{\top},\gamma_{ABE}) &\le \sum_{k_A,k_B} \frac{p_{k_A,k_B}}{1-p^{\perp}} D(|k_A,k_B\rangle\langle k_A,k_B| \otimes \rho_E^{k_A,k_B}, \\ &\quad |k_A,k_A\rangle\langle k_A,k_A| \otimes \rho_E^{k_A,k_B}) = \sum_{k_A\neq k_B} \frac{p_{k_A,k_B}}{1-p^{\perp}} = \frac{1}{1-p^{\perp}} \Pr[K_A\neq K_B]. \end{aligned}$$

Putting this together with Eq. (11), we get

$$\begin{aligned} D(\rho_{ABE},\tilde{\rho}_{ABE}) &= (1-p^{\perp}) D(\rho_{ABE}^{\top}, \tau_{AB}\otimes\rho_E^{\perp}) \\ &\le \Pr[K_A\neq K_B] + (1-p^{\perp}) D(\rho_{AE}^{\top}, \tau_A\otimes\rho_E^{\top}). \end{aligned}$$

■

Proof of Lemma 3.—By construction, $\mathcal{K}_\delta$ aborts with exactly the same probability as the real system. And because σ_E^{QKD} simulates the real protocols, if we plug a converter π_E into $\mathcal{K}\sigma_E^{\mathrm{QKD}}$, which emulates the noisy channel $\mathcal{Q}_q$ and blogs the output of the simulated authentic channel, then $\mathcal{K}_\delta = \mathcal{K}\sigma_E^{\mathrm{QKD}}\pi_E$. Note also that by construction we have $\mathcal{Q}_q\|\mathcal{A}' = (\mathcal{Q}\|\mathcal{A})\pi_E$. Thus,

$$d(\pi_A^{\mathrm{QKD}}\pi_B^{\mathrm{QKD}}(\mathcal{Q}_q\|\mathcal{A}'), \mathcal{K}_\delta) = d(\pi_A^{\mathrm{QKD}}\pi_B^{\mathrm{QKD}}(\mathcal{Q}\|\mathcal{A})\pi_E, \mathcal{K}\sigma_E^{\mathrm{QKD}}\pi_E).$$

Finally, because the converter π_E on both the real and ideal systems can only decrease their distance (see Sec. II.D), the result follows. ■

REFERENCES

Aaronson, Scott, Alexandru Cojocaru, Alexandru Gheorghiu, and Elham Kashefi, 2019, “Complexity-theoretic limitations on blind delegated quantum computation,” in *Proceedings of the 46th International Colloquium on Automata, Languages, and Programming (ICALP 2019), Patras, Greece, 2019*, Leibniz International Proceedings in Informatics Vol. 132, edited by Christel Baier, Ioannis Chatzigiannakis, Paola Flocchini, and Stefano Leonardi (Schloss Dagstuhl, Saarbrücken, Germany), pp. 6:1–6:13.

Acín, Antonio, Nicolas Brunner, Nicolas Gisin, Serge Massar, Stefano Pironio, and Valerio Scarani, 2007, “Device-Independent Security of Quantum Cryptography against Collective Attacks,” Phys. Rev. Lett. **98**, 230501.

Acín, Antonio, Serge Massar, and Stefano Pironio, 2012, “Randomness versus Nonlocality and Entanglement,” Phys. Rev. Lett. **108**, 100402.

Aggarwal, Divesh, Kai-Min Chung, Han-Hsuan Lin, and Thomas Vidick, 2019, “A quantum-proof non-malleable extractor,” in *Advances in Cryptology—EUROCRYPT 2019*, edited by Yuval Ishai and Vincent Rijmen (Springer, New York), pp. 442–469.

Aharonov, Dorit, Michael Ben-Or, and Elad Eban, 2010, “Interactive proofs for quantum computations,” in *Proceedings of Innovations in Computer Science, ICS 2010* (Tsinghua University Press, Beijing), pp. 453–469.

Ahlswede, Rudolph, and Imre Csiszár, 1993, “Common randomness in information theory and cryptography—Part I: Secret sharing,” IEEE Trans. Inf. Theory **39**, 1121–1132.

Alagic, Gorjan, Anne Broadbent, Bill Fefferman, Tommaso Gagliardoni, Christian Schaffner, and Michael St. Jules, 2016, “Computational security of quantum encryption,” in *Proceedings of the 9th International Conference on Information Theoretic Security (ICITS 2016), Tacoma, 2016*, edited by Anderson C. A. Nascimento and Paulo Barreto (Springer, New York), pp. 47–71.

Alagic, Gorjan, Tommaso Gagliardoni, and Christian Majenz, 2018, “Unforgeable quantum encryption,” in *Advances in Cryptology—EUROCRYPT 2018*, Lecture Notes in Computer Science Vol. 10822, edited by Jesper B. Nielsen and Vincent Rijmen (Springer, New York), pp. 489–519.

Alagic, Gorjan, and Christian Majenz, 2017, “Quantum nonmalleability and authentication,” in *Advances in Cryptology—CRYPTO 2017*, Lecture Notes in Computer Science Vol. 10402, edited by Jonathan Katz and Hovav Shacham (Springer, New York), pp. 310–341.

Alicki, Robert, and Mark Fannes, 2004, “Continuity of quantum conditional information,” J. Phys. A **37**, L55–L57.

Alon, Bar, Hao Chung, Kai-Min Chung, Mi-Ying Huang, Yi Lee, and Yu-Ching Shen, 2021, “Round efficient secure multiparty quantum computation with identifiable abort,” in *Advances in Cryptology—CRYPTO 2021*, Lecture Notes in Computer Science Vol. 12828, edited by Tal Malkin and Chris Peikert (Springer, Cham, Switzerland), pp. 436–466.

Ambainis, Andris, Jan Bouda, and Andreas Winter, 2009, “Nonmalleable encryption of quantum information,” J. Math. Phys. (N.Y.) **50**, 042106.

Ambainis, Andris, Michele Mosca, Alain Tapp, and Ronald de Wolf, 2000, “Private quantum channels,” in *Proceedings of the 41st Symposium on Foundations of Computer Science (FOCS ’00), Redondo Beach, CA, 2000* (IEEE, New York), p. 547.

Ambainis, Andris, and Adam Smith, 2004, “Small pseudo-random families of matrices: Derandomizing approximate quantum encryption,” in *Proceedings of the 8th International Workshop on*

Randomization and Computation (RANDOM 2004), Cambridge, MA, 2004 (Springer, New York), pp. 249–260.

Arnon-Friedman, Rotem, 2018, "Reductions to IID in device-independent quantum information processing," Ph.D. thesis (Swiss Federal Institute of Technology Zurich).

Arnon-Friedman, Rotem, Frédéric Dupuis, Omar Fawzi, Renato Renner, and Thomas Vidick, 2018, "Practical device-independent quantum cryptography via entropy accumulation," Nat. Commun. **9**, 459.

Arnon-Friedman, Rotem, Christopher Portmann, and Volkher B. Scholz, 2016, "Quantum-proof multi-source randomness extractors in the Markov model," in *Proceedings of the 11th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2016), Berlin, 2016*, Leibniz International Proceedings in Informatics Vol. 61, edited by Anne Broadbent (Schloss Dagstuhl, Saarbrücken, Germany), pp. 2:1–2:34.

Arnon-Friedman, Rotem, Renato Renner, and Thomas Vidick, 2019, "Simple and tight device-independent security proofs," SIAM J. Comput. **48**, 181–225.

Aspect, Alain, Jean Dalibard, and Gérard Roger, 1982, "Experimental Test of Bell's Inequalities Using Time-Varying Analyzers," Phys. Rev. Lett. **49**, 1804–1807.

Aspect, Alain, Philippe Grangier, and Gérard Roger, 1981, "Experimental Tests of Realistic Local Theories via Bell's Theorem," Phys. Rev. Lett. **47**, 460–463.

Backes, Michael, Birgit Pfitzmann, and Michael Waidner, 2004, "A general composition theorem for secure reactive systems," in *Theory of Cryptography*, Lecture Notes in Computer Science Vol. 2951, edited by Moni Naor (Springer, New York), pp. 336–354.

Backes, Michael, Birgit Pfitzmann, and Michael Waidner, 2007, "The reactive simulatability (RSIM) framework for asynchronous systems," Inf. Comput. **205**, 1685–1720.

Badertscher, Christian, Alexandru Cojocaru, Léo Colisson, Elham Kashefi, Dominik Leichtle, Atul Mantri, and Petros Wallden, 2020, "Security limitations of classical-client delegated quantum computing," in *Advances in Cryptology—ASIACRYPT 2020*, Lecture Notes in Computer Science Vol. 12492, edited by Shiho Moriai and Huaxiong Wang (Springer, Cham, Switzerland), pp. 667–696.

Banfi, Fabio, Ueli Maurer, Christopher Portmann, and Jiamin Zhu, 2019, "Composable and finite computational security of quantum message transmission," in *Theory of Cryptography—TCC 2019*, Lecture Notes in Computer Science Vol. 11891, edited by Dennis Hofheinz and Alon Rosen (Springer, New York), pp. 282–311.

Barnum, Howard, Claude Crépeau, Daniel Gottesman, Adam Smith, and Alain Tapp, 2002, "Authentication of quantum messages," in *Proceedings of the 43rd Symposium on Foundations of Computer Science (FOCS '02), Vancouver, British Columbia, Canada, 2002* (IEEE, New York), pp. 449–458.

Barrett, Jonathan, Roger Colbeck, and Adrian Kent, 2013, "Memory Attacks on Device-Independent Quantum Cryptography," Phys. Rev. Lett. **110**, 010503.

Barrett, Jonathan, Lucien Hardy, and Adrian Kent, 2005, "No Signaling and Quantum Key Distribution," Phys. Rev. Lett. **95**, 010503.

Baumgratz, Tillmann, Marcus Cramer, and Martin B. Plenio, 2014, "Quantifying Coherence," Phys. Rev. Lett. **113**, 140401.

Beaver, Donald, 1992, "Foundations of secure interactive computing," in *Advances in Cryptology—CRYPTO '91*, Lecture Notes in Computer Science Vol. 576, edited by Joan Feigenbaum (Springer, New York), pp. 377–391.

Bell, John Stewart, 1964, "On the Einstein-Podolsky-Rosen paradox," Physics **1**, 195–200.

Bell, John Stewart, 1966, "On the problem of hidden variables in quantum mechanics," Rev. Mod. Phys. **38**, 447–452.

Bell, John Stewart, and Alain Aspect, 2004, "Free variables and local causality," in *Speakable and Unspeakable in Quantum Mechanics: Collected Papers on Quantum Philosophy*, 2nd ed. (Cambridge University Press, Cambridge, England), Chap. 12, pp. 100–104.

Bellare, Mihir, Anand Desai, Eron Jokipii, and Phillip Rogaway, 1997, "A concrete security treatment of symmetric encryption," in *Proceedings of the 38th Annual Symposium on Foundations of Computer Science (FOCS '97), Miami Beach, 1997* (IEEE, New York), pp. 394–403.

Bellare, Mihir, Anand Desai, David Pointcheval, and Phillip Rogaway, 1998, "Relations among notions of security for public-key encryption schemes," in *Advances in Cryptology—CRYPTO '98* (Springer, New York), pp. 26–45.

Bellare, Mihir, and Phillip Rogaway, 2006, "The security of triple encryption and a framework for code-based game-playing proofs," in *Advances in Cryptology—EUROCRYPT 2006*, Lecture Notes in Computer Science Vol. 4004, edited by Serge Vaudenay (Springer, New York), pp. 409–426.

Ben-Aroya, Avraham, and Amnon Ta-Shma, 2012, "Better short-seed quantum-proof extractors, Theor. Comput. Sci. **419**, 17–25.

Bennett, Charles H., Herbert J. Bernstein, Sandu Popescu, and Benjamin Schumacher, 1996, "Concentrating partial entanglement by local operations," Phys. Rev. A **53**, 2046–2052.

Bennett, Charles H., François Bessette, Gilles Brassard, Louis Salvail, and John Smolin, 1992, "Experimental quantum cryptography," J. Cryptol. **5**, 3–28.

Bennett, Charles H., and Gilles Brassard, 1984, "Quantum cryptography: Public key distribution and coin tossing," in *Proceedings of the IEEE International Conference on Computers, Systems, and Signal Processing, Bangalore, India, 1984* (IEEE, New York), pp. 175–179.

Bennett, Charles H., Gilles Brassard, Claude Crépeau, and Ueli Maurer, 1995, "Generalized privacy amplification," IEEE Trans. Inf. Theory **41**, 1915–1923.

Bennett, Charles H., Gilles Brassard, Claude Crépeau, and Marie-Hélène Skubiszewska, 1992, "Practical quantum oblivious transfer," in *Advances in Cryptology—CRYPTO '91*, Lecture Notes in Computer Science Vol. 576 (Springer, New York), pp. 351–366.

Bennett, Charles H., Gilles Brassard, and N. David Mermin, 1992, "Quantum Cryptography without Bell's Theorem," Phys. Rev. Lett. **68**, 557–559.

Bennett, Charles H., Gilles Brassard, Sandu Popescu, Benjamin Schumacher, John A. Smolin, and William K. Wootters, 1996, "Purification of Noisy Entanglement and Faithful Teleportation via Noisy Channels," Phys. Rev. Lett. **76**, 722–725.

Bennett, Charles H., Gilles Brassard, and Jean-Marc Robert, 1988, "Privacy amplification by public discussion," SIAM J. Comput. **17**, 210–229.

Ben-Or, Michael, Claude Crépeau, Daniel Gottesman, Avinatan Hassidim, and Adam Smith, 2006, "Secure multiparty quantum computation with (only) a strict honest majority," in *Proceedings of the 47th Symposium on Foundations of Computer Science (FOCS '06), Berkeley, CA, 2006* (IEEE, New York), pp. 249–260.

Ben-Or, Michael, Michał Horodecki, Debbie Leung, Dominic Mayers, and Jonathan Oppenheim, 2005, "The universal composable security of quantum key distribution," in *Theory of Cryptography—TCC 2005*, Lecture Notes in Computer Science Vol. 3378, edited by Joe Kilian (Springer, New York), pp. 386–406.

Ben-Or, Michael, and Dominic Mayers, 2004, "General security definition and composability for quantum & classical protocols," arXiv:quant-ph/0409062.

Berta, Mario, Matthias Christandl, Roger Colbeck, Joseph M. Renes, and Renato Renner, 2010, "The uncertainty principle in the presence of quantum memory," Nat. Phys. **6**, 659–662.

Berta, Mario, Omar Fawzi, and Volkher B. Scholz, 2017, "Quantum-proof randomness extractors via operator space theory," IEEE Trans. Inf. Theory **63**, 2480–2503.

Biham, Eli, Michel Boyer, Gilles Brassard, Jeroen van de Graaf, and Tal Mor, 2002, "Security of quantum key distribution against all collective attacks," Algorithmica **34**, 372–388.

Biham, Eli, Michel Boyer, P. Oscar Boykin, Tal Mor, and Vwani Roychowdhury, 2000, "A proof of the security of quantum key distribution (extended abstract)," in *Proceedings of the 32nd ACM Symposium on Theory of Computing (STOC '00), Portland, OR, 2000* (ACM, New York), pp. 715–724.

Biham, Eli, Michel Boyer, P. Oscar Boykin, Tal Mor, and Vwani Roychowdhury, 2006, "A proof of the security of quantum key distribution," J. Cryptol. **19**, 381–439.

Biham, Eli, and Tal Mor, 1997, "Security of Quantum Cryptography against Collective Attacks," Phys. Rev. Lett. **78**, 2256–2259.

Blum, Manuel, 1983, "Coin flipping by telephone a protocol for solving impossible problems," ACM SIGACT News **15**, 23–27.

Boileau, J.-C., Kiyoshi Tamaki, Jamie Batuwantudawe, Raymond Laflamme, and Joseph M. Renes, 2005, "Unconditional Security of a Three State Quantum Key Distribution Protocol," Phys. Rev. Lett. **94**, 040503.

Born, Max, 1926, "Zur quantenmechanik der stoßvorgänge," Z. Phys. **37**, 863–867.

Boykin, P. Oscar, and Vwani Roychowdhury, 2003, "Optimal encryption of quantum bits," Phys. Rev. A **67**, 042317.

Bozzio, Mathieu, Eleni Diamanti, and Frédéric Grosshans, 2019, "Semi-device-independent quantum money with coherent states," Phys. Rev. A **99**, 022336.

Branciard, Cyril, Eric G. Cavalcanti, Stephen P. Walborn, Valerio Scarani, and Howard M. Wiseman, 2012, "One-sided device-independent quantum key distribution: Security, feasibility, and the connection with steering," Phys. Rev. A **85**, 010301.

Brandão, Fernando G. S. L., Ravishankar Ramanathan, Andrzej Grudka, Karol Horodecki, Michał Horodecki, Paweł Horodecki, Tomasz Szarek, and Hanna Wojewódka, 2016, "Realistic noise-tolerant randomness amplification using finite number of devices," Nat. Commun. **7**, 11345.

Brassard, Gilles, Claude Crépeau, Dominic Mayers, and Louis Salvail, 1998, "Defeating classical bit commitments with a quantum computer," arXiv:quant-ph/9806031.

Brassard, Gilles, Norbert Lütkenhaus, Tal Mor, and Barry C. Sanders, 2000, "Limitations on Practical Quantum Cryptography," Phys. Rev. Lett. **85**, 1330–1333.

Braunstein, Samuel L., and Stefano Pirandola, 2012, "Side-Channel-Free Quantum Key Distribution," Phys. Rev. Lett. **108**, 130502.

Broadbent, Anne, Joseph Fitzsimons, and Elham Kashefi, 2009, "Universal blind quantum computation," in, *Proceedings of the 50th Symposium on Foundations of Computer Science (FOCS '09), Atlanta, 2009* (IEEE, New York), pp. 517–526.

Broadbent, Anne, Gus Gutoski, and Douglas Stebila, 2013, "Quantum one-time programs," in *Advances in Cryptology—CRYPTO 2013*, Lecture Notes in Computer Science Vol. 8043, edited by Ran Canetti and Juan A. Garay (Springer, New York), pp. 344–360.

Broadbent, Anne, and Stacey Jeffery, 2015, "Quantum homomorphic encryption for circuits of low T-gate complexity," in *Advances in Cryptology—CRYPTO 2015*, edited by Rosario Gennaro and Matthew Robshaw (Springer, New York), pp. 609–629.

Broadbent, Anne, and Christian Schaffner, 2016, "Quantum cryptography beyond quantum key distribution," Des. Codes Cryptogr. **78**, 351–382.

Broadbent, Anne, and Evelyn Wainewright, 2016, "Efficient simulation for quantum message authentication," in *Proceedings of the 9th International Conference on Information Theoretic Security (ICITS 2016), Tacoma, 2016* (Springer, New York), pp. 72–91.

Brunner, Nicolas, Daniel Cavalcanti, Stefano Pironio, Valerio Scarani, and Stephanie Wehner, 2014, "Bell nonlocality," Rev. Mod. Phys. **86**, 419–478.

Buhrman, Harry, Nishanth Chandran, Serge Fehr, Ran Gelles, Vipul Goyal, Rafail Ostrovsky, and Christian Schaffner, 2014, "Position-based quantum cryptography: Impossibility and constructions," SIAM J. Comput. **43**, 150–178.

Calderbank, A. R., and Peter W. Shor, 1996, "Good quantum error-correcting codes exist," Phys. Rev. A **54**, 1098–1105.

Canetti, Ran, 2000, "Security and composition of multiparty cryptographic protocols," J. Cryptol. **13**, 143–202.

Canetti, Ran, 2001, "Universally composable security: A new paradigm for cryptographic protocols," in *Proceedings of the 42nd Symposium on Foundations of Computer Science (FOCS '01), Las Vegas, 2001* (IEEE, New York), pp. 136–145.

Canetti, Ran, 2020, "Universally composable security: A new paradigm for cryptographic protocols," https://eprint.iacr.org/2000/067.pdf.

Canetti, Ran, Ling Cheung, Dilsun Kirli Kaynar, Moses Liskov, Nancy A. Lynch, Olivier Pereira, and Roberto Segala, 2006a, "Task-structured probabilistic I/O automata," in *Proceedings of the 8th International Workshop on Discrete Event Systems (WODES 2006), Ann Arbor, MI, 2006* (IEEE, New York), pp. 207–214, http://theory.csail.mit.edu/~lcheung/papers/task-PIOA-TR.pdf.

Canetti, Ran, Ling Cheung, Dilsun Kirli Kaynar, Moses Liskov, Nancy A. Lynch, Olivier Pereira, and Roberto Segala, 2006b, "Time-bounded task-PIOAs: A framework for analyzing security protocols," in *Proceedings of the 20th International Symposium on Distributed Computing (DISC 2006), Bangalore, India, 2006* (ACM, New York), pp. 238–253.

Canetti, Ran, Yevgeniy Dodis, Rafael Pass, and Shabsi Walfish, 2007, "Universally composable security with global setup," in *Theory of Cryptography—TCC 2007*, Lecture Notes in Computer Science Vol. 4392, edited by Salil P. Vadhan (Springer, New York), pp. 61–85.

Canetti, Ran, and Marc Fischlin, 2001, "Universally composable commitments," in *Advances in Cryptology—CRYPTO 2001*, edited by Joe Kilian (Springer, New York), pp. 19–40.

Canetti, Ran, Hugo Krawczyk, and Jesper B. Nielsen, 2003, "Relaxing chosen-ciphertext security," in *Advances in Cryptology—CRYPTO 2003*, edited by Dan Boneh (Springer, New York), pp. 565–582.

Canetti, Ran, Yehuda Lindell, Rafail Ostrovsky, and Amit Sahai, 2002, "Universally composable two-party and multi-party secure computation," in *Proceedings of the 34th Symposium on Theory of Computing (STOC '02), Montreal, 2002* (ACM, New York), pp. 494–503.

Carter, Larry, and Mark N. Wegman, 1979, "Universal classes of hash functions," J. Comput. Syst. Sci. **18**, 143–154.

Chandran, Nishanth, Vipul Goyal, Ryan Moriarty, and Rafail Ostrovsky, 2009, "Position based cryptography," in *Advances in Cryptology—CRYPTO 2009*, edited by Shai Halevi (Springer, New York), pp. 391–407.

Chen, Yi-Hsiu, Kai-Min Chung, Ching-Yi Lai, Salil P. Vadhan, and Xiaodi Wu, 2017, "Computational notions of quantum min-entropy," arXiv:1704.07309.

Childs, Andrew M., 2005, "Secure assisted quantum computation," Quantum Inf. Comput. **5**, 456–466.

Chiribella, Giulio, Giacomo Mauro D'Ariano, and Paolo Perinotti, 2009, "Theoretical framework for quantum networks," Phys. Rev. A **80**, 022339.

Chitambar, Eric, and Gilad Gour, 2019, "Quantum resource theories," Rev. Mod. Phys. **91**, 025001.

Christandl, Matthias, Artur Ekert, Michał Horodecki, Paweł Horodecki, Jonathan Oppenheim, and Renato Renner, 2007, "Unifying classical and quantum key distillation," in *Theory of Cryptography—TCC 2007*, Lecture Notes in Computer Science Vol. 4392, edited by Salil P. Vadhan (Springer, New York), pp. 456–478.

Christandl, Matthias, Robert König, and Renato Renner, 2009, "Postselection Technique for Quantum Channels with Applications to Quantum Cryptography," Phys. Rev. Lett. **102**, 020504.

Christandl, Matthias, Renato Renner, and Artur Ekert, 2004, "A generic security proof for quantum key distribution," arXiv: quant-ph/0402131.

Christensen, Bradley G., *et al.*, 2013, "Detection-Loophole-Free Test of Quantum Nonlocality, and Applications," Phys. Rev. Lett. **111**, 130406.

Chung, Kai-Min, Xin Li, and Xiaodi Wu, 2014, "Multi-source randomness extractors against quantum side information, and their applications," arXiv:1411.2315.

Chung, Kai-Min, Yaoyun Shi, and Xiaodi Wu, 2014, "Physical randomness extractors: Generating random numbers with minimal assumptions," arXiv:1402.4797.

Clauser, John, Michael Horne, Abner Shimony, and Richard Holt, 1969, "Proposed Experiment to Test Local Hidden-Variable Theories," Phys. Rev. Lett. **23**, 880–884.

Coffman, Valerie, Joydip Kundu, and William K. Wootters, 2000, "Distributed entanglement," Phys. Rev. A **61**, 052306.

Colbeck, Roger, 2006, "Quantum and relativistic protocols for secure multi-party computation," Ph.D. thesis (University of Cambridge).

Colbeck, Roger, and Renato Renner, 2011, "No extension of quantum theory can have improved predictive power," Nat. Commun. **2**, 411.

Colbeck, Roger, and Renato Renner, 2012, "Free randomness can be amplified," Nat. Phys. **8**, 450–454.

Coles, Patrick J., Mario Berta, Marco Tomamichel, and Stephanie Wehner, 2017, "Entropic uncertainty relations and their applications," Rev. Mod. Phys. **89**, 015002.

Conway, John, and Simon Kochen, 2006, "The free will theorem," Found. Phys. **36**, 1441–1473.

Coretti, Sandro, Ueli Maurer, and Björn Tackmann, 2013, "Constructing confidential channels from authenticated channels—Public-key encryption revisited," in *Advances in Cryptology—ASIACRYPT 2013*, Lecture Notes in Computer Science Vol. 8269 (Springer, New York), pp. 134–153.

Cover, Thomas M., and Joy A. Thomas, 2012, *Elements of Information Theory* (John Wiley & Sons, New York).

Cramer, Ronald, Ivan B. Damgård, and Jesper B. Nielsen, 2015, *Secure Multiparty Computation and Secret Sharing* (Cambridge University Press, Cambridge, England).

Crépeau, Claude, Daniel Gottesman, and Adam Smith, 2002, "Secure multi-party quantum computation," in, *Proceedings of the 34th Symposium on Theory of Computing (STOC '02), Montreal, 2002* (ACM, New York), pp. 643–652.

Crépeau, Claude, and Joe Kilian, 1988, "Achieving oblivious transfer using weakened security assumptions," in *Proceedings of the 29th Symposium on Foundations of Computer Science (FOCS '88), White Plains, NY* (ACM, New York), pp. 42–52.

Curty, Marcos, Feihu Xu, Wei Cui, Charles Ci Wen Lim, Kiyoshi Tamaki, and Hoi-Kwong Lo, 2014, "Finite-key analysis for measurement-device-independent quantum key distribution," Nat. Commun. **5**, 3732.

Damgård, Ivan B., Serge Fehr, Louis Salvail, and Christian Schaffner, 2007, "Secure identification and QKD in the bounded-quantum-storage model," in *Advances in Cryptology—CRYPTO 2007*, edited by Alfred Menezes (Springer, New York), pp. 342–359.

Damgård, Ivan B., Serge Fehr, Louis Salvail, and Christian Schaffner, 2008, "Cryptography in the bounded-quantum-storage model," SIAM J. Comput. **37**, 1865–1890.

De, Anindya, Christopher Portmann, Thomas Vidick, and Renato Renner, 2012, "Trevisan's extractor in the presence of quantum side information," SIAM J. Comput. **41**, 915–940.

Demay, Gregory, and Ueli Maurer, 2013, "Unfair coin tossing," in *Proceedings of the 2013 IEEE International Symposium on Information Theory (ISIT 2013), Istanbul, 2013* (IEEE, New York), pp. 1556–1560.

Devetak, Igor, and Andreas Winter, 2005, "Distillation of secret key and entanglement from quantum states," Proc. R. Soc. A **461**, 207–235.

Dickinson, Paul, and Ashwin Nayak, 2006, "Approximate randomization of quantum states with fewer bits of key," AIP Conf. Proc. **864**, 18–36.

DiVincenzo, David, Michał Horodecki, Debbie Leung, John Smolin, and Barbara Terhal, 2004, "Locking Classical Correlation in Quantum States," Phys. Rev. Lett. **92**, 067902.

Dodis, Yevgeniy, and Daniel Wichs, 2009, "Non-malleable extractors and symmetric key cryptography from weak secrets," in *Proceedings of the 41st Symposium on Theory of Computing (STOC '09), Bethesda, MD, 2009* (ACM, New York), pp. 601–610.

Dulek, Yfke, Alex B. Grilo, Stacey Jeffery, Christian Majenz, and Christian Schaffner, 2020, "Secure multi-party quantum computation with a dishonest majority," in *Advances in Cryptology—EUROCRYPT 2020*, edited by Anne Canteaut and Yuval Ishai (Springer, New York), pp. 729–758.

Dunjko, Vedran, Joseph Fitzsimons, Christopher Portmann, and Renato Renner, 2014, "Composable security of delegated quantum computation," in *Advances in Cryptology—ASIACRYPT 2014*, Lecture Notes in Computer Science Vol. 8874, edited by Palash Sarkar and Tetsu Iwata (Springer, New York), pp. 406–425.

Dunjko, Vedran, and Elham Kashefi, 2016, "Blind quantum computing with two almost identical states, arXiv:1604.01586.

Dupuis, Frédéric, Omar Fawzi, and Renato Renner, 2020, "Entropy accumulation," Commun. Math. Phys. **379**, 867–913.

Dupuis, Frédéric, Jesper B. Nielsen, and Louis Salvail, 2012, "Actively secure two-party evaluation of any quantum operation," in *Advances in Cryptology—CRYPTO 2012*, Lecture Notes in Computer Science Vol. 7417, edited by Reihaneh Safavi-Naini and Ran Canetti (Springer, New York), pp. 794–811.

Einstein, Albert, Boris Podolsky, and Nathan Rosen, 1935, "Can quantum-mechanical description of physical reality be considered complete?," Phys. Rev. **47**, 777–780.

Ekert, Artur, 1991, "Quantum Cryptography Based on Bell's Theorem," Phys. Rev. Lett. **67**, 661–663.

Ekert, Artur, and Renato Renner, 2014, "The ultimate physical limits of privacy," Nature (London) **507**, 443–447.

Elkouss, David, Anthony Leverrier, Romain Alleaume, and Joseph J. Boutros, 2009, "Efficient reconciliation protocol for discrete-variable quantum key distribution," in *Proceedings of the 2009 IEEE International Symposium on Information Theory (ISIT 2009), Seoul, 2009* (IEEE, New York), pp. 1879–1883.

Elkouss, David, Jesus Martinez-Mateo, and Vicente Martin, 2011, "Information reconciliation for quantum key distribution," Quantum Inf. Comput. **11**, 226–238.

Fehr, Serge, and Christian Schaffner, 2008, "Randomness extraction via δ-biased masking in the presence of a quantum attacker," in *Theory of Cryptography—TCC 2008*, Lecture Notes in Computer Science Vol. 4948, edited by Ran Canetti (Springer, New York), pp. 465–481.

Fitzsimons, Joseph F., and Elham Kashefi, 2017, "Unconditionally verifiable blind computation," Phys. Rev. A **96**, 012303.

Freedman, Stuart J., and John F. Clauser, 1972, "Experimental Test of Local Hidden-Variable Theories," Phys. Rev. Lett. **28**, 938–941.

Fuchs, Christopher A., 1998, "Information gain vs. state disturbance in quantum theory," Fortschr. Phys. **46**, 535–565.

Fuchs, Christopher A., Nicolas Gisin, Robert B. Griffiths, Chi-Sheng Niu, and Asher Peres, 1997, "Optimal eavesdropping in quantum cryptography. I. Information bound and optimal strategy," Phys. Rev. A **56**, 1163–1172.

Fuchs, Christopher A., and Jeroen Van De Graaf, 1999, "Cryptographic distinguishability measures for quantum-mechanical states," IEEE Trans. Inf. Theory **45**, 1216–1227.

Fung, Chi-Hang Fred, Bing Qi, Kiyoshi Tamaki, and Hoi-Kwong Lo, 2007, "Phase-remapping attack in practical quantum-key-distribution systems," Phys. Rev. A **75**, 032314.

Garg, Sumegha, Henry Yuen, and Mark Zhandry, 2017, "New security notions and feasibility results for authentication of quantum data," in *Advances in Cryptology—CRYPTO 2017*, Lecture Notes in Computer Science Vol. 10402, edited by Jonathan Katz and Hovav Shacham (Springer, New York), pp. 342–371.

Gavinsky, Dmitry, Julia Kempe, Iordanis Kerenidis, Ran Raz, and Ronald de Wolf, 2007, "Exponential separations for one-way quantum communication complexity, with applications to cryptography," in *Proceedings of the 39th Symposium on Theory of Computing (STOC '07), San Diego, 2007* (ACM, New York), pp. 516–525.

Gerhardt, Ilja, Qin Liu, Antía Lamas-Linares, Johannes Skaar, Christian Kurtsiefer, and Vadim Makarov, 2011, "Full-field implementation of a perfect eavesdropper on a quantum cryptography system," Nat. Commun. **2**, 349.

Gheorghiu, Alexandru, and Thomas Vidick, 2019, "Computationally-secure and composable remote state preparation," in *Proceedings of the IEEE 60th Symposium on Foundations of Computer Science (FOCS '19), Baltimore, 2019* (IEEE, New York), pp. 1024–1033.

Gisin, Nicolas, Sylvain Fasel, Barbara Kraus, Hugo Zbinden, and Grégoire Ribordy, 2006, "Trojan-horse attacks on quantum-key-distribution systems," Phys. Rev. A **73**, 022320.

Giustina, Marissa, *et al.*, 2013, "Bell violation using entangled photons without the fair-sampling assumption," Nature (London) **497**, 227–230.

Giustina, Marissa, *et al.*, 2015, "Significant-Loophole-Free Test of Bell's Theorem with Entangled Photons," Phys. Rev. Lett. **115**, 250401.

Goldreich, Oded, 2004, *Foundations of Cryptography: Vol. 2—Basic Applications* (Cambridge University Press, New York).

Goldreich, Oded, Silvia Micali, and Avi Wigderson, 1987, "How to play any mental game," in," *Proceedings of the 19th Symposium on Theory of Computing (STOC '87), New York, 1987* (ACM, New York), pp. 218–229.

Goldreich, Oded, Silvio Micali, and Avi Wigderson, 1986, "Proofs that yield nothing but their validity and a methodology of cryptographic protocol design," in *Proceedings of the 27th Symposium on Foundations of Computer Science (FOCS '86), Toronto, 1986* (IEEE, New York), pp. 174–187.

Gottesman, Daniel, and Hoi-Kwong Lo, 2003, "Proof of security of quantum key distribution with two-way classical communications," IEEE Trans. Inf. Theory **49**, 457–475.

Gottesman, Daniel, Hoi-Kwong Lo, Norbert Lütkenhaus, and John Preskill, 2004, "Security of quantum key distribution with imperfect devices," Quantum Inf. Comput. **4**, 325–360.

Goyal, Vipul, Yuval Ishai, Amit Sahai, Ramarathnam Venkatesan, and Akshay Wadia, 2010, "Founding cryptography on tamper-proof hardware tokens," in *Theory of Cryptography—TCC 2010*, Lecture Notes in Computer Science Vol. 5978, edited by Daniele Micciancio (Springer, New York), pp. 308–326.

Gutoski, Gus, 2012, "On a measure of distance for quantum strategies," J. Math. Phys. (N.Y.) **53**, 032202.

Gutoski, Gus, and John Watrous, 2007, "Toward a general theory of quantum games," in *Proceedings of the 39th Symposium on Theory of Computing (STOC '07), San Diego, 2007* (ACM, New York), pp. 565–574.

Hardy, Lucien, 2005, "Probability theories with dynamic causal structure: A new framework for quantum gravity," arXiv:gr-qc/0509120.

Hardy, Lucien, 2007, "Towards quantum gravity: A framework for probabilistic theories with non-fixed causal structure," J. Phys. A **40**, 3081.

Hardy, Lucien, 2011, "Reformulating and reconstructing quantum theory," arXiv:1104.2066.

Hardy, Lucien, 2012, "The operator tensor formulation of quantum theory," Phil. Trans. R. Soc. A **370**, 3385–3417.

Hardy, Lucien, 2015, "Quantum theory with bold operator tensors," Phil. Trans. R. Soc. A **373**, 20140239.

Hayashi, Masahito, and Toyohiro Tsurumaru, 2012, "Concise and tight security analysis of the Bennett-Brassard 1984 protocol with finite key lengths, New J. Phys. **14**, 093014.

Hayden, Patrick, Debbie Leung, and Dominic Mayers, 2011, "The universal composable security of quantum message authentication with key recycling," arXiv:1610.09434.

Hayden, Patrick, Debbie Leung, Peter W. Shor, and Andreas Winter, 2004, "Randomizing quantum states: Constructions and applications," Commun. Math. Phys. **250**, 371–391.

Helstrom, Carl W., 1976, *Quantum Detection and Estimation Theory*, Mathematics in Science and Engineering Vol. 123, (Academic Press, New York).

Hensen, B., *et al.*, 2015, "Loophole-free Bell inequality violation using electron spins separated by 1.3 kilometres," Nature (London) **526**, 682–686.

Hofheinz, Dennis, Jörn Müller-Quade, and Dominique Unruh, 2006, "On the (im)possibility of extending coin toss," in *Advances in Cryptology—EUROCRYPT 2006*, Lecture Notes in Computer Science Vol. 4004 (Springer, New York), pp. 504–521.

Hofheinz, Dennis, and Victor Shoup, 2015, "GNUC: A new universal composability framework," J. Cryptol. **28**, 423–508.

Horodecki, Karol, Michał Horodecki, Paweł Horodecki, Debbie Leung, and Jonathan Oppenheim, 2008, "Quantum key distribution based on private states: Unconditional security over untrusted channels with zero quantum capacity," IEEE Trans. Inf. Theory **54**, 2604–2620.

Horodecki, Karol, and Maciej Stankiewicz, 2020, "Semi-device-independent quantum money," New J. Phys. **22**, 023007.

Horodecki, Michał, Paweł Horodecki, and Ryszard Horodecki, 1998, "Mixed-State Entanglement and Distillation: Is There a 'Bound' Entanglement in Nature?," Phys. Rev. Lett. **80**, 5239–5242.

Hwang, Won-Young, 2003, "Quantum Key Distribution with High Loss: Toward Global Secure Communication," Phys. Rev. Lett. **91**, 057901.

Inamori, Hitoshi, Norbert Lütkenhaus, and Dominic Mayers, 2007, "Unconditional security of practical quantum key distribution," Eur. Phys. J. D **41**, 599–627.

Inoue, Kyo, Edo Waks, and Yoshihisa Yamamoto, 2002, "Differential Phase Shift Quantum Key Distribution," Phys. Rev. Lett. **89**, 037902.

Ishai, Yuval, Rafail Ostrovsky, and Vassilis Zikas, 2014, "Secure multi-party computation with identifiable abort," in *Advances in Cryptology—CRYPTO 2014*, edited by Juan A. Garay and Rosario Gennaro (Springer, New York), pp. 369–386.

Ishai, Yuval, Manoj Prabhakaran, and Amit Sahai, 2008, "Founding cryptography on oblivious transfer—efficiently," in *Advances in Cryptology—CRYPTO 2008*, Lecture Notes in Computer Science Vol. 5157, edited by David Wagner (Springer, New York), pp. 572–591.

Jost, Daniel, and Ueli Maurer, 2018, "Security definitions for hash functions: Combining UCE and Indifferentiability," in *International Conference on Security and Cryptography for Networks—SCN 2018*, Lecture Notes in Computer Science Vol. 11035, edited by Dario Catalano and Roberto De Prisco (Springer, New York), pp. 83–101.

Jouguet, Paul, and Sebastien Kunz-Jacques, 2014, "High performance error correction for quantum key distribution using polar codes," Quantum Inf. Comput. **14**, 329–338.

Kaniewski, Jędrzej, 2015, "Relativistic quantum cryptography," Ph.D. thesis (National University of Singapore).

Kaniewski, Jędrzej, Marco Tomamichel, Esther Hänggi, and Stephanie Wehner, 2013, "Secure bit commitment from relativistic constraints," IEEE Trans. Inf. Theory **59**, 4687–4699.

Katz, Jonathan, and Moti Yung, 2006, "Characterization of security notions for probabilistic private-key encryption," J. Cryptol. **19**, 67–95.

Kent, Adrian, 1999, "Unconditionally Secure Bit Commitment," Phys. Rev. Lett. **83**, 1447–1450.

Kent, Adrian, 2012, "Unconditionally Secure Bit Commitment by Transmitting Measurement Outcomes," Phys. Rev. Lett. **109**, 130501.

Kessler, Max, and Rotem Arnon-Friedman, 2020, "Device-independent randomness amplification and privatization," IEEE J. Sel. Areas Inf. Theory **1**, 568–584.

Koashi, Masato, 2004, "Unconditional Security of Coherent-State Quantum Key Distribution with a Strong Phase-Reference Pulse," Phys. Rev. Lett. **93**, 120501.

Koashi, Masato, 2009, "Simple security proof of quantum key distribution based on complementarity," New J. Phys. **11**, 045018.

Koashi, Masato, and Andreas Winter, 2004, "Monogamy of quantum entanglement and other correlations," Phys. Rev. A **69**, 022309.

Kochen, Simon B., and Ernst P. Specker, 1967, "The problem of hidden variables in quantum mechanics," J. Math. Mech. **17**, 59–87, http://www.jstor.org/stable/24902153.

König, Robert, Ueli Maurer, and Renato Renner, 2005, "On the power of quantum memory," IEEE Trans. Inf. Theory **51**, 2391–2401.

König, Robert, Renato Renner, Andor Bariska, and Ueli Maurer, 2007, "Small Accessible Quantum Information Does Not Imply Security," Phys. Rev. Lett. **98**, 140502.

König, Robert, and Barbara M. Terhal, 2008, "The bounded-storage model in the presence of a quantum adversary," IEEE Trans. Inf. Theory **54**, 749–762.

König, Robert, Stephanie Wehner, and Jürg Wullschleger, 2012, "Unconditional security from noisy quantum storage," IEEE Trans. Inf. Theory **58**, 1962–1984.

Kraus, Barbara, Nicolas Gisin, and Renner Renner, 2005, "Lower and Upper Bounds on the Secret-Key Rate for Quantum Key Distribution Protocols Using One-Way Classical Communication," Phys. Rev. Lett. **95**, 080501.

Küsters, Ralf, 2006, "Simulation-based security with inexhaustible interactive turing machines," in *Proceedings of the 19th IEEE workshop on Computer Security Foundations (CSFW '06), Venice, 2006* (IEEE, New York), pp. 309–320.

Laneve, Lorenzo, and Lídia del Rio, 2021, "Impossibility of composable oblivious transfer in relativistic quantum cryptography," arXiv:2106.11200.

Leverrier, Anthony, Romain Alléaume, Joseph Boutros, Gilles Zémor, and Philippe Grangier, 2008, "Multidimensional reconciliation for a continuous-variable quantum key distribution," Phys. Rev. A **77**, 042325.

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

Lim, Charles Ci Wen, Christopher Portmann, Marco Tomamichel, Renato Renner, and Nicolas Gisin, 2013, "Device-Independent Quantum Key Distribution with Local Bell Test," Phys. Rev. X **3**, 031006.

Lipinska, Victoria, Jérémy Ribeiro, and Stephanie Wehner, 2020, "Secure multiparty quantum computation with few qubits," Phys. Rev. A **102**, 022405.

Liu, Yang, *et al.*, 2013, "Experimental Measurement-Device-Independent Quantum Key Distribution," Phys. Rev. Lett. **111**, 130502.

Liu, Yi-Kai, 2014, "Single-shot security for one-time memories in the isolated qubits model," in *Advances in Cryptology—CRYPTO 2014*, edited by Juan A. Garay and Rosario Gennaro (Springer, New York), pp. 19–36.

Liu, Yi-Kai, 2015, "Privacy amplification in the isolated qubits model," in *Advances in Cryptology—EUROCRYPT 2015*, edited by Elisabeth Oswald and Marc Fischlin (Springer, New York), pp. 785–814.

Lo, Hoi-Kwong, 2003, "Method for decoupling error correction from privacy amplification," New J. Phys. **5**, 36–36.

Lo, Hoi-Kwong, and Hoi Fung Chau, 1999, "Unconditional security of quantum key distribution over arbitrarily long distances," Science **283**, 2050–2056.

Lo, Hoi-Kwong, Marcos Curty, and Bing Qi, 2012, "Measurement-Device-Independent Quantum Key Distribution," Phys. Rev. Lett. **108**, 130503.

Lo, Hoi-Kwong, Xiongfeng Ma, and Kai Chen, 2005, "Decoy State Quantum Key Distribution," Phys. Rev. Lett. **94**, 230504.

Lucamarini, Marco, Zhiliang L. Yuan, James F. Dynes, and Andrew J. Shields, 2018, "Overcoming the rate-distance limit of quantum key distribution without quantum repeaters," Nature (London) **557**, 400–403.

Lütkenhaus, Norbert, 2000, "Security against individual attacks for realistic quantum key distribution," Phys. Rev. A **61**, 052304.

Lydersen, Lars, Carlos Wiechers, Christoffer Wittmann, Dominique Elser, Johannes Skaar, and Vadim Makarov, 2010, "Hacking commercial quantum cryptography systems by tailored bright illumination," Nat. Photonics **4**, 686–689.

Ma, Jiajun, You Zhou, Xiao Yuan, and Xiongfeng Ma, 2019, "Operational interpretation of coherence in quantum key distribution," Phys. Rev. A **99**, 062325.

Ma, Xiongfeng, and Mohsen Razavi, 2012, "Alternative schemes for measurement-device-independent quantum key distribution," Phys. Rev. A **86**, 062319.

Makarov, Vadim, 2009, "Controlling passively quenched single photon detectors by bright light," New J. Phys. **11**, 065003.

Makarov, Vadim, Andrey Anisimov, and Johannes Skaar, 2006, "Effects of detector efficiency mismatch on security of quantum cryptosystems," Phys. Rev. A **74**, 022313.

Mateus, Paulo, John C. Mitchell, and Andre Scedrov, 2003, "Composition of cryptographic protocols in a probabilistic polynomial-time process calculus," in *CONCUR 2003—Concurrency Theory*, Lecture Notes in Computer Science Vol. 2761, edited by Yang Xiang and Brahim Chaib-draa (Springer, New York), pp. 327–349.

Mauerer, Wolfgang, Christopher Portmann, and Volkher B. Scholz, 2012, "A modular framework for randomness extraction based on trevisan's construction," arXiv:1212.0520.

Maurer, Ueli, 1993, "Secret key agreement by public discussion," IEEE Trans. Inf. Theory **39**, 733–742.

Maurer, Ueli, 1994, "The strong secret key rate of discrete random triples," in *Communications and Cryptography: Two Sides of One Tapestry*, Springer International Series in Engineering and Computer Science Vol. 276, edited by Richard E. Blahut, Daniel J. Costello, Jr., and Ueli Maurer (Springer, New York), pp. 271–285.

Maurer, Ueli, 2002, "Indistinguishability of random systems," in *Advances in Cryptology—EUROCRYPT 2002*, Lecture Notes in Computer Science Vol. 2332 (Springer, New York), pp. 110–132.

Maurer, Ueli, 2012, "Constructive cryptography—A new paradigm for security definitions and proofs," in *Theory of Security and Applications—TOSCA 2011*, Lecture Notes in Computer Science Vol. 6993, edited by Sebastian Mödersheim and Catuscia Palamidessi (Springer, New York), pp. 33–56.

Maurer, Ueli, Krzysztof Pietrzak, and Renato Renner, 2007, "Indistinguishability amplification," in *Advances in Cryptology—CRYPTO 2007*, Lecture Notes in Computer Science Vol. 4622 (Springer, New York), pp. 130–149.

Maurer, Ueli, and Renato Renner, 2011, "Abstract cryptography," in *Proceedings of Innovations in Computer Science (ICS 2011), Beijing, 2011* (Tsinghua University Press, Beijing), pp. 1–21.

Maurer, Ueli, and Renato Renner, 2016, "From indifferentiability to constructive cryptography (and back)," in *Theory of Cryptography—TCC 2016*, Lecture Notes in Computer Science Vol. 9985, edited by Martin Hirt and Adam Smith (Springer, New York), pp. 3–24.

Maurer, Ueli, Andreas Rüedlinger, and Björn Tackmann, 2012, "Confidentiality and integrity: A constructive perspective," in *Theory of Cryptography—TCC 2012*, Lecture Notes in Computer Science Vol. 7194, edited by Ronald Cramer (Springer, New York), pp. 209–229.

Maurer, Ueli, and Stefan Wolf, 2000, "Information-theoretic key agreement: From weak to strong secrecy for free," in *Advances in Cryptology—EUROCRYPT 2000*, Lecture Notes in Computer Science Vol. 1807, edited by Bart Preneel (Springer, New York), pp. 351–368.

Mayers, Dominic, 1996, "Quantum key distribution and string oblivious transfer in noisy channels," in *Advances in Cryptology—CRYPTO '96*, Lecture Notes in Computer Science Vol. 1109, edited by Neal Koblitz (Springer, New York), pp. 343–357.

Mayers, Dominic, 2001, "Unconditional security in quantum cryptography," J. ACM **48**, 351–406.

Micali, Silvio, and Phillip Rogaway, 1992, "Secure computation (abstract)," in *Advances in Cryptology—CRYPTO '91*, Lecture Notes in Computer Science Vol. 576, edited by Joan Feigenbaum (Springer, New York), pp. 392–404.

Miller, Carl, and Yaoyun Shi, 2014, "Robust protocols for securely expanding randomness and distributing keys using untrusted quantum devices," in *Proceedings of the 46th Symposium on Theory of Computing (STOC '14), New York, 2014* (ACM, New York), pp. 417–426.

Mitchell, John C., Ajith Ramanathan, Andre Scedrov, and Vanessa Teague, 2006, "A probabilistic polynomial-time process calculus for the analysis of cryptographic protocols," Theor. Comput. Sci. **353**, 118–164.

Muller, Antoine, Thomas Herzog, Bruno Huttner, Woflgang Tittel, Hugo Zbinden, and Nicolas Gisin, 1997, " 'Plug and play' systems for quantum cryptography," Appl. Phys. Lett. **70**, 793–795.

Müller-Quade, Jörn, and Renato Renner, 2009, "Composability in quantum cryptography," New J. Phys. **11**, 085006.

Nielsen, Michael A., and Isaac L. Chuang, 2010, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, England).

Nishioka, Tsuyoshi, Hirokazu Ishizuka, Toshio, and Junichi Abe, 2002, " 'Circular type' quantum key distribution," IEEE Photonics Technol. Lett. **14**, 576–578.

Ohya, Masanori, and Dénes Petz, 1993, *Quantum Entropy and Its Use* (Springer, New York).

Pawłowski, Marcin, and Nicolas Brunner, 2011, "Semi-device-independent security of one-way quantum key distribution," Phys. Rev. A **84**, 010302.

Peres, Asher, and Daniel R. Terno, 2004, "Quantum information and relativity theory," Rev. Mod. Phys. **76**, 93–123.

Pfitzmann, Birgit, and Michael Waidner, 2000, "Composition and integrity preservation of secure reactive systems," in *Proceedings of the 7th ACM Conference on Computer and Communications Security (CSS '00), Athens, 2000* (ACM, New York), pp. 245–254.

Pfitzmann, Birgit, and Michael Waidner, 2001, "A model for asynchronous reactive systems and its application to secure message transmission," in *Proceedings of the 2001 IEEE Symposium on Security and Privacy, Oakland, 2001* (IEEE, New York), pp. 184–200.

Pirandola, Stefano, Carlo Ottaviani, Gaetana Spedalieri, Christian Weedbrook, Samuel L. Braunstein, Seth Lloyd, Tobias Gehring, Christian S. Jacobsen, and Ulrik L. Andersen, 2015, "High-rate measurement-device-independent quantum cryptography," Nat. Photonics **9**, 397.

Pironio, Stefano, Antonio Acín, Nicolas Brunner, Nicolas Gisin, Serge Massar, and Valerio Scarani, 2009, "Device-independent quantum key distribution secure against collective attacks," New J. Phys. **11**, 045021.

Pironio, Stefano, Antonio Acín, Serge Massar, A. Boyer de La Giroday, Dzimitry N. Matsukevich, Peter Maunz, Steven Olmschenk, David Hayes, Le Luo, and T. Andrew Manning, 2010, "Random numbers certified by Bell's theorem," Nature (London) **464**, 1021–1024.

Portmann, Christopher, 2014, "Key recycling in authentication," IEEE Trans. Inf. Theory **60**, 4383–4396.

Portmann, Christopher, 2017a, "Quantum authentication with key recycling," in *Advances in Cryptology—EUROCRYPT 2017*, Lecture Notes in Computer Science Vol. 10212 (Springer, New York), pp. 339–368.

Portmann, Christopher, 2017b, "(Quantum) min-entropy resources," arXiv:1705.10595.

Portmann, Christopher, Christian Matt, Ueli Maurer, Renato Renner, and Björn Tackmann, 2017, "Causal boxes: Quantum information-processing systems closed under composition," IEEE Trans. Inf. Theory **63**, 3277–3305.

Prokop, Milŏs, 2020, "Composable security of quantum bit commitment protocol," https://project-archive.inf.ed.ac.uk/ug4/20201685/ug4_proj.pdf.

Qi, Bing, Chi-Hang Fred Fung, Hoi-Kwong Lo, and Xiongfeng Ma, 2007, "Time-shift attack in practical quantum cryptosystems," Quantum Inf. Comput. **7**, 73–82.

Reichardt, Ben W., Falk Unger, and Umesh Vazirani, 2013, "Classical command of quantum systems," Nature (London) **496**, 456–460.

Renes, Joseph M., 2013, "The physics of quantum information: Complementarity, uncertainty, and entanglement," Int. J. Quantum. Inf. **11**, 1330002.

Renes, Joseph M., and Renato Renner, 2012, "One-shot classical data compression with quantum side information and the distillation of common randomness or secret keys," IEEE Trans. Inf. Theory **58**, 1985–1991.

Renes, Joseph M., and Renato Renner, 2020, "Are quantum cryptographic security claims vacuous?," arXiv:2010.11961.

Renner, Renato, 2005, "Security of quantum key distribution," Ph.D. thesis (Swiss Federal Institute of Technology Zurich).

Renner, Renato, 2007, "Symmetry of large physical systems implies independence of subsystems," Nat. Phys. **3**, 645–649.

Renner, Renato, Nicolas Gisin, and Barbara Kraus, 2005, "Information-theoretic security proof for quantum-key-distribution protocols," Phys. Rev. A **72**, 012332.

Renner, Renato, and Robert König, 2005, "Universally composable privacy amplification against quantum adversaries," in *Theory of Cryptography—TCC 2005*, Lecture Notes in Computer Science Vol. 3378, edited by Joe Kilian (Springer, New York), pp. 407–425.

Renner, Renato, and Stefan Wolf, 2003, "Unconditional authenticity and privacy from an arbitrarily weak secret," in *Advances in Cryptology—CRYPTO 2003*, Lecture Notes in Computer Science Vol. 2729, edited by Dan Boneh (Springer, New York), pp. 78–95.

Renner, Renato, and Stefan Wolf, 2005, "Simple and tight bounds for information reconciliation and privacy amplification," in *Advances in Cryptology—ASIACRYPT 2005*, Lecture Notes in Computer Science Vol. 3788, edited by Bimal Roy (Springer, New York), pp. 199–216.

Renner, Renner, and J. Ignacio Cirac, 2009, "de Finetti Representation Theorem for Infinite-Dimensional Quantum Systems and Applications to Quantum Cryptography," Phys. Rev. Lett. **102**, 110504.

Rivest, Ronald L., Adi Shamir, and Leonard Adleman, 1978, "A method for obtaining digital signatures and public-key cryptosystems," Commun. ACM **21**, 120–126.

Rogaway, Phillip, 2006, "Formalizing human ignorance," in *Progress in Cryptology—VIETCRYPT 2006*, Lecture Notes in Computer Science Vol. 4341, edited by Phong Q. Nguyen (Springer, New York), pp. 211–228.

Rosenfeld, Wenjamin, Daniel Burchardt, Robert Garthoff, Kai Redeker, Norbert Ortegel, Markus Rau, and Harald Weinfurter, 2017, "Event-Ready Bell Test Using Entangled Atoms Simultaneously Closing Detection and Locality Loopholes," Phys. Rev. Lett. **119**, 010402.

Rowe, M. A., David Kielpinski, V. Meyer, Charles A. Sackett, Wayne M. Itano, C. Monroe, and D. J. Wineland, 2001, "Experimental violation of a Bell's inequality with efficient detection," Nature (London) **409**, 791–794.

Sasaki, Toshihiko, Yoshihisa Yamamoto, and Masato Koashi, 2014, "Practical quantum key distribution protocol without monitoring signal disturbance," Nature (London) **509**, 475.

Scarani, Valerio, 2013, "The device-independent outlook on quantum physics (lecture notes on the power of Bell's theorem)," arXiv:1303.3081.

Scarani, Valerio, Antonio Acín, Grégoire Ribordy, and Nicolas Gisin, 2004, "Quantum Cryptography Protocols Robust against Photon Number Splitting Attacks for Weak Laser Pulse Implementations," Phys. Rev. Lett. **92**, 057901.

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

Scarani, Valerio, and Renato Renner, 2008, "Quantum Cryptography with Finite Resources: Unconditional Security Bound for Discrete-Variable Protocols with One-Way Postprocessing," Phys. Rev. Lett. **100**, 200501.

Schaffner, Christian, Barbara Terhal, and Stephanie Wehner, 2009, "Robust cryptography in the noisy-quantum-storage model," Quantum Inf. Comput. **9**, 963–996.

Seiler, Gregor, and Ueli Maurer, 2016, "On the impossibility of information-theoretic composable coin toss extension," in *Proceedings of the 2016 IEEE International Symposium on Information Theory (ISIT 2016), Barcelona, 2016* (IEEE, New York), pp. 3058–3061.

Shalm, Lynden K., *et al.*, 2015, "Strong Loophole-Free Test of Local Realism," Phys. Rev. Lett. **115**, 250402.

Shaltiel, Ronen, 2004, "Recent developments in explicit constructions of extractors," in *Current Trends in Theoretical Computer Science: The Challenge of the New Century, Vol. 1—Algorithms and Complexity*, edited by Gheorghe Plun, Grzegorz Rozenberg, and Arto Salomaa (World Scientific, Singapore), pp. 189–228.

Shannon, Claude E., 1949, "Communication theory of secrecy systems," Bell Syst. Tech. J. **28**, 656–715.

Sheridan, Lana, Phuc Le Thinh, and Valerio Scarani, 2010, "Finite-key security against coherent attacks in quantum key distribution," New J. Phys. **12**, 123019.

Shor, Peter W., 1997, "Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer," SIAM J. Comput. **26**, 1484–1509.

Shor, Peter W., and John Preskill, 2000, "Simple Proof of Security of the BB84 Quantum Key Distribution Protocol," Phys. Rev. Lett. **85**, 441–444.

Simmons, Gustavus J., 1985, "Authentication theory/coding theory," in *Advances in Cryptology—CRYPTO '84*, Lecture Notes in Computer Science Vol. 196, edited by George Robert Blakley and David Chaum (Springer, New York), pp. 411–431.

Simmons, Gustavus J., 1988, "A survey of information authentication," Proc. IEEE **76**, 603–620.

Steane, Andrew, 1996, "Multiple-particle interference and quantum error correction," Proc. R. Soc. A **452**, 2551–2577.

Stinson, Douglas R., 1990, "The combinatorics of authentication and secrecy codes," J. Cryptol. **2**, 23–49.

Stinson, Douglas R., 1994, "Universal hashing and authentication codes," Des. Codes Cryptogr. **4**, 369–380.

Streltsov, Alexander, Gerardo Adesso, and Martin B. Plenio, 2017, "Colloquium: Quantum coherence as a resource," Rev. Mod. Phys. **89**, 041003.

Stucki, Damien, Nicolas Brunner, Nicolas Gisin, Valerio Scarani, and Hugo Zbinden, 2005, "Fast and simple one-way quantum key distribution," Appl. Phys. Lett. **87**, 194108.

Tamaki, Kiyoshi, Masato Koashi, and Nobuyuki Imoto, 2003, "Unconditionally Secure Key Distribution Based on Two Nonorthogonal States," Phys. Rev. Lett. **90**, 167904.

Tamaki, Kiyoshi, and Hoi-Kwong Lo, 2006, "Unconditionally secure key distillation from multiphotons," Phys. Rev. A **73**, 010302.

Tamaki, Kiyoshi, Hoi-Kwong Lo, Chi-Hang Fred Fung, and Bing Qi, 2012, "Phase encoding schemes for measurement-device-independent quantum key distribution with basis-dependent flaw," Phys. Rev. A **85**, 042307.

Tan, Ernest Y.-Z., Charles Ci Wen Lim, and Renato Renner, 2020, "Advantage Distillation for Device-Independent Quantum Key Distribution," Phys. Rev. Lett. **124**, 020502.

Tang, Yan-Lin, *et al.*, 2014, "Measurement-Device-Independent Quantum Key Distribution over 200 km," Phys. Rev. Lett. **113**, 190501.

Terhal, Barbara M., 2004, "Is entanglement monogamous?," IBM J. Res. Dev. **48**, 71–78.

Thorisson, Hermann, 2000, *Coupling, Stationarity, and Regeneration*, Probability and Its Applications (Springer, New York).

Tittel, Wolfgang, Jurgen Brendel, Bernard Gisin, Thomas Herzog, Hugo Zbinden, and Nicolas Gisin, 1998, "Experimental demonstration of quantum correlations over more than 10 km," Phys. Rev. A **57**, 3229–3232.

Tomamichel, Marco, and Anthony Leverrier, 2017, "A largely self-contained and complete security proof for quantum key distribution," Quantum **1**, 14.

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

Tomamichel, Marco, and Renato Renner, 2011, "Uncertainty Relation for Smooth Entropies," Phys. Rev. Lett. **106**, 110506.

Tomamichel, Marco, Christian Schaffner, Adam Smith, and Renato Renner, 2010, "Leftover hashing against quantum side information," in *Proceedings of the 2010 IEEE International Symposium on Information Theory (ISIT 2010), Austin, 2010* (IEEE, New York), pp. 2703–2707.

Trushechkin, Anton S., 2020, "On the operational meaning and practical aspects of using the security parameter in quantum key distribution," Quantum Electron. **50**, 426–439.

Unruh, Dominique, 2004, "Simulatable security for quantum protocols," arXiv:quant-ph/0409125.

Unruh, Dominique, 2010, "Universally composable quantum multi-party computation," in *Advances in Cryptology—EUROCRYPT 2010*, Lecture Notes in Computer Science Vol. 6110, edited by Henri Gilbert (Springer, New York), pp. 486–505.

Unruh, Dominique, 2011, "Concurrent composition in the bounded quantum storage model," in *Advances in Cryptology—EUROCRYPT 2011*, Lecture Notes in Computer Science Vol. 6632, edited by Kenneth G. Paterson (Springer, New York), pp. 467–486.

Unruh, Dominique, 2013, "Everlasting multi-party computation," in *Advances in Cryptology—CRYPTO 2013*, Lecture Notes in Computer Science Vol. 8043, edited by Ran Canetti and Juan A. Garay (Springer, New York), pp. 380–397.

Unruh, Dominique, 2014, "Quantum position verification in the random oracle model," in *Advances in Cryptology—CRYPTO 2014*, Lecture Notes in Computer Science Vol. 8617, edited by Juan A. Garay and Rosario Gennaro (Springer, New York), pp. 1–18.

Vakhitov, Artem, Vadim Makarov, and Dag R. Hjelme, 2001, "Large pulse attack as a method of conventional optical eavesdropping in quantum cryptography," J. Mod. Opt. **48**, 2023–2038.

Vazirani, Umesh, and Thomas Vidick, 2012, "Certifiable quantum dice: Or, true random number generation secure against quantum adversaries," in *Proceedings of the 44th Symposium on Theory of Computing (STOC '12), New York, 2012* (ACM, New York), pp. 61–76.

Vazirani, Umesh, and Thomas Vidick, 2014, "Fully Device-Independent Quantum Key Distribution," Phys. Rev. Lett. **113**, 140501.

Vernam, Gilbert S., 1926, "Cipher printing telegraph systems for secret wire and radio telegraphic communications," Trans. Am. Inst. Electr. Eng. **XLV**, 295–301.

Vilasini, V., Christopher Portmann, and Lídia del Rio, 2019, "Composable security in relativistic quantum cryptography," New J. Phys. **21**, 043057.

Wang, Xiang-Bin, 2005, "Beating the Photon-Number-Splitting Attack in Practical Quantum Cryptography," Phys. Rev. Lett. **94**, 230503.

Watrous, John, 2018, *The Theory of Quantum Information* (Cambridge University Press, Cambridge, England), http://cs.uwaterloo.ca/~watrous/TQI/.

Webb, Zak, 2016, "The Clifford group forms a unitary 3-design," Quantum Inf. Comput. **16**, 1379–1400.

Wegman, Mark N., and Larry Carter, 1981, "New hash functions and their use in authentication and set equality," J. Comput. Syst. Sci. **22**, 265–279.

Wehner, Stephanie, Christian Schaffner, and Barbara M. Terhal, 2008, "Cryptography from Noisy Storage," Phys. Rev. Lett. **100**, 220502.

Weier, Henning, Harald Krauss, Markus Rau, Martin Fürst, Sebastian Nauerth, and Harald Weinfurter, 2011, "Quantum eavesdropping without interception: An attack exploiting the dead time of single-photon detectors," New J. Phys. **13**, 073024.

Weihs, Gregor, Thomas Jennewein, Christoph Simon, Harald Weinfurter, and Anton Zeilinger, 1998, "Violation of Bell's Inequality Under Strict Einstein Locality Conditions," Phys. Rev. Lett. **81**, 5039–5043.

Wiesner, Stephen, 1983, "Conjugate coding," ACM SIGACT News **15**, 78–88.

Winter, Andreas, 2017, "Weak locking capacity of quantum channels can be much larger than private capacity," J. Cryptol. **30**, 1–21.

Wolf, Stefan, 1999, "Information-theoretically and computationally secure key agreement in cryptography," Ph.D. thesis (Swiss Federal Institute of Technology Zurich).

Wootters, William K., and Wojciech H. Zurek, 1982, "A single quantum cannot be cloned," Nature (London) **299**, 802–803.

Xu, Feihu, Bing Qi, and Hoi-Kwong Lo, 2010, "Experimental demonstration of phase-remapping attack in a practical quantum key distribution system," New J. Phys. **12**, 113026.

Yin, Hua-Lei, and Zeng-Bing Chen, 2019, "Finite-key analysis for twin-field quantum key distribution with composable security," Sci. Rep. **9**, 17113.

Yin, Hua-Lei, *et al.*, 2016, "Measurement-Device-Independent Quantum Key Distribution over a 404 km Optical Fiber," Phys. Rev. Lett. **117**, 190501.

Yuan, Zhiliang, James F. Dynes, and Andrew J. Shields, 2010, "Avoiding the blinding attack in QKD," Nat. Photonics **4**, 800.

Zhandry, Mark, 2012, "How to construct quantum random functions," in *Proceedings of the 53rd Symposium on Foundations of Computer Science (FOCS '12), New Brunswick, NJ, 2012* (IEEE, New York), pp. 679–687.

Zhao, Yi, Chi-Hang Fred Fung, Bing Qi, Christine Chen, and Hoi-Kwong Lo, 2008, "Quantum hacking: Experimental demonstration of time-shift attack against practical quantum-key-distribution systems," Phys. Rev. A **78**, 042333.

Zhu, Huangjun, 2017, "Multiqubit Clifford groups are unitary 3-designs," Phys. Rev. A **96**, 062336.

Zuckerman, David, 1990, "General weak random sources," in *Proceedings of the 31st Symposium on Foundations of Computer Science (FOCS '90), St. Louis, 1990* (IEEE, New York), pp. 534–543.