

Revisiting Deniability in Quantum Key Exchange via Covert Communication and Entanglement Distillation

Arash Atashpendar^{1(*)}, G. Vamsi Policharla², Peter B. Rønne¹, Peter Y.A. Ryan¹

¹ SnT, University of Luxembourg, Luxembourg

{arash.atashpendar, peter.roenne, peter.ryan}@uni.lu

² Department of Physics, Indian Institute of Technology Bombay, India
guruvamsi.policharla@iitb.ac.in

Abstract. We revisit the notion of deniability in quantum key exchange (QKE), a topic that remains largely unexplored. In the only work on this subject by Donald Beaver, it is argued that QKE is not necessarily deniable due to an eavesdropping attack that limits key equivocation. We provide more insight into the nature of this attack and how it extends to other constructions such as QKE obtained from uncloneable encryption. We then adopt the framework for quantum authenticated key exchange, developed by Mosca et al., and extend it to introduce the notion of coercer-deniable QKE, formalized in terms of the indistinguishability of real and fake coercer views. Next, we apply results from a recent work by Arrazola and Scarani on covert quantum communication to establish a connection between covert QKE and deniability. We propose DC-QKE, a simple deniable covert QKE protocol, and prove its deniability via a reduction to the security of covert QKE. Finally, we consider how entanglement distillation can be used to enable information-theoretically deniable protocols for QKE and tasks beyond key exchange.

1 Introduction

Deniability represents a fundamental privacy-related notion in cryptography. The ability to deny a message or an action is a desired property in many contexts such as off-the-record communication, anonymous reporting, whistle-blowing and coercion-resistant secure electronic voting. The concept of non-repudiation is closely related to deniability in that the former is aimed at associating specific actions with legitimate parties and thereby preventing them from denying that they have performed a certain task, whereas the latter achieves the opposite property by allowing legitimate parties to deny having performed a particular action. For this reason, deniability is sometimes referred to as *repudiability*.

The definitions and requirements for deniable exchange can vary depending on the cryptographic task in question, e.g., encryption, authentication or key exchange. Roughly speaking, the common underlying idea for a deniable scheme can be understood as the impossibility for an adversary to produce cryptographic proofs, using only algorithmic evidence, that would allow a third-party, often referred to as a judge, to decide if a particular entity has either taken part in a given exchange or exchanged

a certain message, which can be a secret key, a digital signature, or a plaintext message. In the context of key exchange, this can be also formulated in terms of a corrupt party (receiver) proving to a judge that a message can be traced back to the other party [16].

In the public-key setting, an immediate challenge for achieving deniability is posed by the need for remote authentication as it typically gives rise to binding evidence, e.g., digital signatures, see [16,17]. The formal analysis of deniability in classical cryptography can be traced back to the original works of Canetti et al. and Dwork et al. on deniable encryption [11] and deniable authentication [18], respectively. These led to a series of papers on this topic covering a relatively wide array of applications. Deniable key exchange was first formalized by Di Raimondo et al. in [16] using a framework based on the simulation paradigm, which is closely related to that of zero-knowledge proofs.

Despite being a well-known and fundamental concept in classical cryptography, rather surprisingly, deniability has been largely ignored by the quantum cryptography community. To put things into perspective, with the exception of a single paper by Donald Beaver [3], and a footnote in [20] commenting on the former, there are no other works that directly tackle deniable QKE.

In the adversarial setting described in [3], it is assumed that the honest parties are approached by the adversary after the termination of a QKE session and demanded to reveal their private randomness, i.e., the raw key bits encoded in their quantum states. It is then claimed that QKE schemes, despite having perfect and unconditional security, are not necessarily deniable due to an eavesdropping attack. In the case of the BB84 protocol, this attack introduces a binding between the parties' inputs and the final key, thus constraining the space of the final secret key such that key equivocation is no longer possible.

Note that since Beaver's work [3] appeared a few years before a formal analysis of deniability for key exchange was published, its analysis is partly based on the adversarial model formulated earlier in [11] for deniable encryption. For this reason, the setting corresponds more closely to scenarios wherein the honest parties try to deceive a coercer by presenting fake messages and randomness, e.g., deceiving a coercer who tries to verify a voter's claimed choice using an intercepted ciphertext of a ballot in the context of secure e-voting.

1.1 Contributions and Structure

In Section 3 we revisit the notion of deniability in QKE and provide more insight into the eavesdropping attack aimed at detecting attempts at denial described in [3]. Having shed light on the nature of this attack, we show that while coercer-deniability can be achieved by uncloneable encryption (UE) [19], QKE obtained from UE remains vulnerable to the same attack. We briefly elaborate on the differences between our model and simulation-based deniability [16]. To provide a firm foundation, we adopt the framework and security model for quantum authenticated key exchange (Q-AKE) developed by Mosca et al. [24] and extend them to introduce the notion of coercer-deniable QKE, which we formalize in terms of the indistinguishability of real and fake coercer views.

We establish a connection between the concept of covert communication and deniability in Section 4, which to the best of our knowledge has not been formally considered before. More precisely, we apply results from a recent work by Arrazola and Scarani on obtaining covert quantum communication and covert QKE via noise injection [1] to propose DC-QKE, a simple construction for coercer-deniable QKE. We prove the deniability of DC-QKE via a reduction to the security of covert QKE. Compared to the candidate PQECC protocol suggested in [3] that is claimed to be deniable, our construction does not require quantum computation and falls within the more practical realm of prepare-and-measure protocols.

Finally, in Section 5 we consider how quantum entanglement distillation can be used not only to counter eavesdropping attacks, but also to achieve information-theoretic deniability. We conclude by presenting some open questions in Section 6. It is our hope that this work will rekindle interest, more broadly, in the notion of deniable communication in the quantum setting, a topic that has received very little attention from the quantum cryptography community.

1.2 Related Work

We focus on some of the most prominent works in the extensive body of work on deniability in classical cryptography. The notion of deniable encryption was considered by Canetti et al. [11] in a setting where an adversary demands that parties reveal private coins used for generating a ciphertext. This motivated the need for schemes equipped with a faking algorithm that can produce fake randomness with distributions indistinguishable from that of the real encryption.

In a framework based on the simulation paradigm, Dwork et al. introduced the notion of deniable authentication [18], followed by the work of Di Raimondo et al. on the formalization of deniable key exchange [16]. Both works rely on the formalism of zero-knowledge (ZK) proofs, with definitions formalized in terms of a simulator that can produce a simulated view that is indistinguishable from the real one. In a subsequent work, Di Raimondo and Gennaro gave a formal definition of forward deniability [15], requiring that indistinguishability remain intact even when a (corrupted) party reveals real coins after a session. Among other things, they showed that statistical ZK protocols are forward deniable.

Pass [26] formally defines the notion of deniable zero-knowledge and presents positive and negative results in the common reference string and random oracle model. In [17], Dodis et al. establish a link between deniability and ideal authentication and further model a situation in which deniability should hold even when a corrupted party colludes with the adversary during the execution of a protocol. They show an impossibility result in the PKI model if adaptive corruptions are allowed. Cremers and Feltz introduced another variant for key exchange referred to as peer and time deniability [13], while also capturing perfect forward secrecy. More recently, Unger and Goldberg studied deniable authenticated key exchange (DAKE) in the context of secure messaging [31].

To the best of our knowledge, the only work related to deniability in QKE is a single paper by Beaver [3], in which the author suggests a negative result arguing that existing QKE schemes are not necessarily deniable.

2 Preliminaries in Quantum Information and QKE

We use the Dirac bra-ket notation and standard terminology from quantum computing. Here we limit ourselves to a description of the most relevant concepts in quantum information theory. More details can be found in standard textbooks [25,32]. For brevity, let A and B denote the honest parties, and E the adversary.

Given an orthonormal basis formed by $|0\rangle$ and $|1\rangle$ in a two-dimensional complex Hilbert space \mathcal{H}_2 , let $(+)$ $\equiv \{|0\rangle, |1\rangle\}$ denote the computational basis and (\times) $\equiv \{(1/\sqrt{2})(|0\rangle + |1\rangle), (1/\sqrt{2})(|0\rangle - |1\rangle)\}$ the diagonal basis.

If the state vector of a composite system cannot be expressed as a tensor product $|\psi_1\rangle \otimes |\psi_2\rangle$, the state of each subsystem cannot be described independently and we say the two qubits are *entangled*. This property is best exemplified by maximally entangled qubits (*ebits*), the so-called *Bell states*

$$|\Phi^\pm\rangle_{AB} = \frac{1}{\sqrt{2}}(|00\rangle_{AB} \pm |11\rangle_{AB}) \quad , \quad |\Psi^\pm\rangle_{AB} = \frac{1}{\sqrt{2}}(|01\rangle_{AB} \pm |10\rangle_{AB})$$

A noisy qubit that cannot be expressed as a linear superposition of pure states is said to be in a *mixed* state, a classical probability distribution of pure states: $\{p_X(x), |\psi_x\rangle\}_{x \in X}$. The *density operator* ρ , defined as a weighted sum of projectors, captures both pure and mixed states: $\rho \equiv \sum_{x \in X} p_X(x) |\psi_x\rangle \langle \psi_x|$.

Given a density matrix ρ_{AB} describing the joint state of a system held by A and B , the *partial trace* allows us to compute the local state of A (density operator ρ_A) if B 's system is not accessible to A . To obtain ρ_A from ρ_{AB} (the reduced state of ρ_{AB} on A), we trace out the system B : $\rho_A = \text{Tr}_B(\rho_{AB})$. As a distance measure, we use the expected fidelity $F(|\psi\rangle, \rho)$ between a pure state $|\psi\rangle$ and a mixed state ρ given by $F(|\psi\rangle, \rho) = \langle \psi | \rho | \psi \rangle$.

A crucial distinction between quantum and classical information is captured by the well-known No-Cloning theorem [33], which states that an arbitrary unknown quantum state cannot be copied or cloned perfectly.

2.1 Quantum Key Exchange and Uncloneable Encryption

QKE allows two parties to establish a common secret key with information-theoretic security using an insecure quantum channel, and a public authenticated classical channel. In Protocol 1 we describe the **BB84** protocol, the most well-known QKE variant due to Bennett and Brassard [5]. For consistency with related works, we use the well-established formalism based on error-correcting codes, developed by Shor and Preskill [28]. Let $C_1[n, k_1]$ and $C_2[n, k_2]$ be two classical linear binary codes encoding k_1 and k_2 bits in n bits such that $\{0\} \subset C_2 \subset C_1 \subset \mathbf{F}_2^n$ where \mathbf{F}_2^n is the binary vector space on n bits. A mapping of vectors $v \in C_1$ to a set of basis states (codewords) for the Calderbank-Shor-Steane (CSS) [10,29] code subspace is given by: $v \mapsto (1/\sqrt{|C_2|}) \sum_{w \in C_2} |v + w\rangle$. Due to the irrelevance of phase errors and their decoupling from bit flips in CSS codes, Alice can send $|v\rangle$ along with classical error-correction information $u + v$ where $u, v \in \mathbf{F}_2^n$ and $u \in C_1$, such that Bob can decode to a codeword in C_1 from $(v + \epsilon) - (u + v)$ where ϵ is an error codeword, with the final key being the coset leader of $u + C_2$.

Protocol 1 BB84 for an n -bit key with protection against δn bit errors

- 1: Alice generates two random bit strings $a, b \in \{0, 1\}^{(4+\delta)n}$, encodes a_i into $|\psi_i\rangle$ in basis (+) if $b_i = 0$ and in (\times) otherwise, and $\forall i \in [1, |a|]$ sends $|\psi_i\rangle$ to Bob.
 - 2: Bob generates a random bit string $b' \in \{0, 1\}^{(4+\delta)n}$ and upon receiving the qubits, measures $|\psi_i\rangle$ in (+) or (\times) according to b'_i to obtain a'_i .
 - 3: Alice announces b and Bob discards a'_i where $b_i \neq b'_i$, ending up with at least $2n$ bits with high probability.
 - 4: Alice picks a set p of $2n$ bits at random from a , and a set q containing n elements of p chosen as check bits at random. Let $v = p \setminus q$.
 - 5: Alice and Bob compare their check bits and abort if the error exceeds a predefined threshold.
 - 6: Alice announces $u + v$, where v is the string of the remaining non-check bits, and u is a random codeword in C_1 .
 - 7: Bob subtracts $u + v$ from his code qubits, $v + \epsilon$, and corrects the result, $u + \epsilon$, to a codeword in C_1 .
 - 8: Alice and Bob use the coset of $u + C_2$ as their final secret key of length n .
-

Uncloneable encryption (UE) enables transmission of ciphertexts that cannot be perfectly copied and stored for later decoding, by encoding carefully prepared codewords into quantum states, thereby leveraging the No-Cloning theorem. We refer to Gottesman's original work [19] for a detailed explanation of the sketch in Protocol 2. Alice and Bob agree on a message length n , a Message Authentication Code (MAC) of length s , an error-correcting code C_1 having message length K and codeword length N with distance $2\delta N$ for average error rate δ , and another error-correcting code C_2 (for privacy amplification) with message length K' and codeword length N and distance $2(\delta + \eta)N$ to correct more errors than C_1 , satisfying $C_2^\perp \subset C_1$, where C_2^\perp is the dual code containing all vectors orthogonal to C_2 . The pre-shared key is broken down into four pieces, all chosen uniformly at random: an authentication key $k \in \{0, 1\}^s$, a one-time pad $e \in \{0, 1\}^{n+s}$, a syndrome $c_1 \in \{0, 1\}^{N-K}$, and a basis sequence $b \in \{0, 1\}^N$.

Protocol 2 Uncloneable Encryption for sending a message $m \in \{0, 1\}^n$

- 1: Compute $\text{MAC}(m)_k = \mu \in \{0, 1\}^s$. Let $x = m || \mu \in \{0, 1\}^{n+s}$.
- 2: Mask x with the one-time pad e to obtain $y = x \oplus e$.
- 3: From the coset of C_1 given by the syndrome c_1 , pick a random codeword $z \in \{0, 1\}^N$ that has syndrome bits y w.r.t. C_2^\perp , where $C_2^\perp \subset C_1$.
- 4: For $i \in [1, N]$ encode ciphertext bit z_i in the basis (+) if $b_i = 0$ and in the basis (\times) if $b_i = 1$. The resulting state $|\psi_i\rangle$ is sent to Bob.

To perform decryption:

- 1: For $i \in [1, N]$, measure $|\psi_i\rangle$ according to b_i , to obtain $z'_i \in \{0, 1\}^N$.
 - 2: Perform error-correction on z' using code C_1 and evaluate the parity checks of C_2/C_1^\perp for privacy amplification to get an $(n + s)$ -bit string y' .
 - 3: Invert the OTP step to obtain $x' = y' \oplus e$.
 - 4: Parse x' as the concatenation $m' || \mu'$ and use k to verify if $\text{MAC}(m')_k = \mu'$.
-

QKE from UE. It is known [19] that any quantum authentication (QA) scheme can be used as a secure UE scheme, which can in turn be used to obtain QKE, with less interaction and more efficient error detection. We give a brief description of how QKE can be obtained from UE in Protocol 3.

Protocol 3 Obtaining QKE from Uncloneable Encryption

- 1: Alice generates random strings k and x , and sends x to Bob via UE, keyed with k .
 - 2: Bob announces that he has received the message, and then Alice announces k .
 - 3: Bob decodes the classical message x , and upon MAC verification, if the message is valid, he announces this to Alice and they will use x as their secret key.
-

3 Coercer-Deniable Quantum Key Exchange

Following the setting in [3], in which it is implicitly assumed that the adversary has established a binding between the participants' identities and a given QKE session, we introduce the notion of coercer-deniability for QKE. This makes it possible to consider an adversarial setting similar to that of deniable encryption [11] and expect that the parties might be coerced into revealing their private coins after the termination of a session, in which case they would have to produce fake randomness such that the resulting transcript and the claimed values remain consistent with the adversary's observations.

Beaver's analysis [3] is briefly addressed in a footnote in a paper by Ioannou and Mosca [20] and the issue is brushed aside based on the argument that the parties do not have to keep records of their raw key bits. It is argued that for deniability to be satisfied, it is sufficient that the adversary cannot provide binding evidence that attributes a particular key to the classical communication as their measurements on the quantum channel do not constitute a publicly verifiable proof. However, counter-arguments for this view were already raised in the motivations for deniable encryption [11] in terms of secure erasure being difficult and unreliable, and that erasing cannot be externally verified. Moreover, it is also argued that if one were to make the physical security assumption that random choices made for encryption are physically unavailable, the deniability problem would disappear. We refer to [11] and references therein for more details.

Bindings, or lack thereof, lie at the core of deniability. Although we leave a formal comparison of our model with the one formulated in the simulation paradigm [16] as future work, a notable difference can be expressed in terms of the inputs presented to the adversary. In the simulation paradigm, deniability is modelled only according to the simulatability of the legal transcript that the adversary or a corrupt party produces naturally via a session with a party as evidence for the judge, whereas for coercer-deniability, the adversary additionally demands that the honest parties reveal their private randomness.

Finally, note that viewing deniability in terms of "convincing" the adversary is bound to be problematic and indeed a source of debate in the cryptographic research

community as the adversary may never be convinced given their knowledge of the existence of faking algorithms. Hence, deniability is formulated in terms of the indistinguishability of views (or their simulatability [16]) such that a judge would have no reason to believe a given transcript provided by the adversary establishes a binding as it could have been forged or simulated.

3.1 Defeating Deniability in QKE via Eavesdropping in a Nutshell

We briefly review the eavesdropping attack described in [3] and provide further insight. Suppose Alice sends qubit $|\psi\rangle^{m,b}$ to Bob, which encodes a single-bit message m prepared in a basis determined by $b \in \{+, \times\}$. Let $\Phi(E, m)$ denote the state obtained after sending $|\psi\rangle^{m,b}$, relayed and possibly modified by an adversary E . Moreover, let $\rho(E, m)$ denote the view presented to the judge, obtained by tracing over inaccessible systems. Now for a qubit measured correctly by Eve, if a party tries to deny by pretending to have sent $\sigma_1 = \rho(E, 1)$ instead of $\sigma_2 = \rho(E, 0)$, e.g., by using some local transformation U_{neg} to simply negate a given qubit, then $F(\sigma_1, \sigma_2) = 0$, where F denotes the fidelity between σ_1 and σ_2 . Thus, the judge can successfully detect this attempt at denial.

This attack can be mounted successfully with non-negligible probability without causing the session to abort: Assume that N qubits will be transmitted in a BB84 session and that the tolerable error rate is $\frac{\eta}{N}$, where clearly $\eta \sim N$. Eve measures each qubit with probability $\frac{\eta}{N}$ (choosing a basis at random) and passes on the remaining ones to Bob undisturbed, i.e., she plants a number of decoy states proportional to the tolerated error threshold. On average, $\frac{\eta}{2}$ measurements will come from matching bases, which can be used by Eve to detect attempts at denial, if Alice claims to have measured a different encoding. After discarding half the qubits in the sifting phase, this ratio will remain unchanged. Now Alice and/or Bob must flip at least one bit in order to deny without knowledge of where the decoy states lie in the transmitted sequence, thus getting caught with probability $\frac{\eta}{2N}$ upon flipping a bit at random.

3.2 On the Coercer-Deniability of Uncloneable Encryption

The vulnerability described in Section 3.1 is made possible by an eavesdropping attack that induces a binding in the key coming from a BB84 session. Uncloneable encryption remains immune to this attack because the quantum encoding is done for an already one-time padded classical input. More precisely, a binding established at the level of quantum states can still be perfectly denied because the actual raw information bits m are not directly encoded into the sequence of qubits, instead the concatenation of m and the corresponding authentication tag $\mu = \text{MAC}_k(m)$, i.e., $x = m||\mu$, is masked with a one-time pad e to obtain $y = x \oplus e$, which is then mapped onto a codeword z that is encoded into quantum states. For this reason, in the context of coercer-deniability, regardless of a binding established on z by the adversary, Alice can still deny to another input message in that she can pick a different input $x' = m'||\mu'$ to compute a fake pad $e' = y \oplus x'$, so that upon revealing e' to Eve, she will simply decode $y \oplus e' = x'$, as intended.

However, note that a prepare-and-measure QKE obtained from UE still remains vulnerable to the same eavesdropping attack due to the fact that we can no longer make use of the deniability of the one-time pad in UE such that the bindings induced by Eve constrain the choice of the underlying codewords.

3.3 Security Model

We adopt the framework for quantum AKEs developed by Mosca et al. [24]. Due to space constraints, we mainly focus on our proposed extensions. **Parties**, including the adversary, are modelled as a pair of classical and quantum Turing machines (TM) that execute a series of interactive computations and exchange messages with each other through classical and quantum channels, collectively referred to as a **protocol**. An execution of a protocol is referred to as a **session**, identified with a unique session identifier. An ongoing session is called an *active* session, and upon completion, it either outputs an error term \perp in case of an abort, or it outputs a tuple $(sk, pid, \mathbf{v}, \mathbf{u})$ in case of a successful termination. The tuple consists of a session key sk , a party identifier pid and two vectors \mathbf{u} and \mathbf{v} that model public values and secret terms, respectively.

We adopt an extended version of the **adversarial model** described in [24], to account for coercer-deniability. Let E be an efficient, i.e. (quantum) polynomial time, adversary with classical and quantum runtime bounds $t_c(k)$ and $t_q(k)$, and quantum memory bound $m_q(k)$, where bounds can be unlimited. Following standard assumptions, the adversary controls all communication between parties and carries the messages exchanged between them. We consider an authenticated classical channel and do not impose any special restrictions otherwise. Additionally, the adversary is allowed to approach either the sender or the receiver after the termination of a session and request access to a subset $\mathbf{r} \subseteq \mathbf{v}$ of the private randomness used by the parties for a given session, i.e. set of values to be faked.

Security notions can be formulated in terms of **security experiments** in which the adversary interacts with the parties via a set of well-defined **queries**. These queries typically involve sending messages to an active session or initiating one, corrupting a party, learning their long-term secret key, revealing the ephemeral keys of an incomplete session, obtaining the computed session key for a given session, and a **test-session**(id) query capturing the winning condition of the game that can be invoked only for a *fresh* session. Revealing secret values to the adversary is modeled via **partnering**. The notion of *freshness* captures the idea of excluding cases that would allow the adversary to trivially win the security experiment. This is done by imposing minimal restrictions on the set of queries the adversary can invoke for a given session such that there exist protocols that can still satisfy the definition of session-key security. A session remains fresh as long as at least one element in \mathbf{u} and \mathbf{v} remains secret, see [24] for more details.

The **transcript** of a protocol consists of all publicly exchanged messages between the parties during a run or session of the protocol. The definition of “views” and “outputs” given in [3] coincides with that of transcripts in [16] in the sense that it allows us to model a transcript that can be obtained from observations made on the quantum channel. The *view* of a party P consists of their state in \mathcal{H}_P along with any

classical strings they produce or observe. More generally, for a two-party protocol, captured by the global density matrix ρ_{AB} for the systems of A and B , the individual system A corresponds to a partial trace that yields a reduced density matrix, i.e., $\rho_A = \text{Tr}_B(\rho_{AB})$, with a similar approach for any additional couplings.

3.4 Coercer-Deniable QKE via View Indistinguishability

We use the security model in Section 3.3 to introduce the notion of coercer-deniable QKE, formalized via the indistinguishability of real and fake views. Note that in this work we do not account for forward deniability and forward secrecy.

Coercer-Deniability Security Experiment. Let $\text{CoercerDenQKE}_{E,C}^{\Pi}(\kappa)$ denote this experiment and Q the same set of queries available to the adversary in a security game for session-key security, as described in Section 3.3, and [24]. Clearly, in addition to deniability, it is vital that the security of the session key remains intact as well. For this reason, we simply extend the requirements of the security game for a session-key secure KE by having the challenger \mathcal{C} provide an additional piece of information to the adversary E when the latter calls the `test-session()` query. This means that the definition of a fresh session remains the same as the one given in [24]. E invokes queries from $Q \setminus \{\text{test-session}()\}$ until E issues `test-session()` to a fresh session of their choice. \mathcal{C} decides on a random bit b and if $b = 0$, \mathcal{C} provides E with the real session key k and the real vector of private randomness \mathbf{r} , and if $b = 1$, with a random (fake) key k' and a random (fake) vector of private randomness \mathbf{r}' . Finally, E guesses an output b' and wins the game if $b = b'$. The experiment returns 1 if E succeeds, and 0 otherwise. Let $\text{Adv}_E^{\Pi}(\kappa) = |\Pr[b = b'] - 1/2|$ denote the winning advantage of E .

Definition 1 (Coercer-Deniable QKE). For adversary E , let there be an efficient distinguisher D_E on security parameter κ . We say that $\Pi_{\mathbf{r}}$ is a coercer-deniable QKE protocol if, for any adversary E , transcript \mathbf{t} , and for any k, k' , and a vector of private random inputs $\mathbf{r} = (r_1, \dots, r_\ell)$, there exists a denial/faking program $\mathcal{F}_{A,B}$ that running on $(k, k', \mathbf{t}, \mathbf{r})$ produces $\mathbf{r}' = (r'_1, \dots, r'_\ell)$ such that the following conditions hold:

- Π is a secure QKE protocol.
- The adversary E cannot do better than making a random guess for winning the coercer-deniability security experiment, i.e., $\text{Adv}_E^{\Pi}(\kappa) \leq \text{negl}(\kappa)$

$$\Pr[\text{CoercerDenQKE}_{E,C}^{\Pi}(\kappa) = 1] \leq \frac{1}{2} + \text{negl}(\kappa)$$

Equivalently, we require that for all efficient distinguisher D_E

$$|\Pr[D_E(\text{View}_{\text{Real}}(k, \mathbf{t}, \mathbf{r})) = 1] - \Pr[D_E(\text{View}_{\text{Fake}}(k', \mathbf{t}, \mathbf{r}')) = 1]| \leq \text{negl}(\kappa),$$

where the transcript $\mathbf{t} = (\mathbf{c}, \rho_E(k))$ is a tuple consisting of a vector \mathbf{c} , containing classical message exchanges of a session, along with the local view of the adversary w.r.t. the quantum channel obtained by tracing over inaccessible systems (see Section 3.3).

A function $f : \mathbb{N} \rightarrow \mathbb{R}$ is negligible if for any constant k , there exists a N_k such that $\forall N \geq N_k$, we have $f(N) < N^{-k}$. In other words, it approaches zero faster than any polynomial in the asymptotic limit.

Remark 1. We introduced a vector of private random inputs \mathbf{r} to avoid being restricted to a specific set of “fake coins” in a coercer-deniable setting such as the raw key bits in BB84 as used in Beaver’s analysis. This allows us to include other private inputs as part of the transcript that need to be forged by the denying parties without having to provide a new security model for each variant. Indeed, in [24], Mosca et al. consider the security of QKE in case various secret values are compromised before or after a session. This means that these values can, in principle, be included in the set of random coins that might have to be revealed to the adversary and it should therefore be possible to generate fake alternatives using a faking algorithm.

4 Deniable QKE via Covert Quantum Communication

We establish a connection between covert communication and deniability by providing a simple construction for coercer-deniable QKE using covert QKE. We then show that deniability is reduced to the covertness property, meaning that deniable QKE can be performed as long as covert QKE is not broken by the adversary, formalized via the security reduction given in Theorem 2.

Covert communication becomes relevant when parties wish to keep the very act of communicating secret or hidden from a malicious warden. This can be motivated by various requirements such as the need for hiding one’s communication with a particular entity when this act alone can be incriminating. While encryption can make it impossible for the adversary to access the contents of a message, it would not prevent them from detecting exchanges over a channel under their observation. Bash et al. [2,27] established a square-root law for covert communication in the presence of an unbounded quantum adversary stating that $\mathcal{O}(\sqrt{n})$ covert bits can be exchanged over n channel uses. Recently, Arrazola and Scarani [1] extended covert communication to the quantum regime for transmitting qubits covertly. Covert quantum communication consists of two parties exchanging a sequence of qubits such that an adversary trying to detect this cannot succeed by doing better than making a random guess, i.e., $P_d \leq \frac{1}{2} + \epsilon$ for sufficiently small $\epsilon > 0$, where P_d denotes the probability of detection and ϵ the detection bias.

4.1 Covert Quantum Key Exchange

Since covert communication requires pre-shared secret randomness, a natural question to ask is whether QKE can be done covertly. This was also addressed in [1] and it was shown that covert QKE with unconditional security for the covertness property is impossible because the amount of key consumption is greater than the amount produced. However, a hybrid approach involving pseudo-random number generators (PRNG) was proposed to achieve covert QKE with a positive key rate such that the resulting secret key remains information-theoretically secure, while the covertness of

QKE is shown to be at least as strong as the security of the PRNG. The PRNG is used to expand a truly random pre-shared key into an exponentially larger pseudo-random output, which is then used to determine the time-bins for sending signals in covert QKE.

Covert QKE Security Experiment. Let $\text{CovertQKE}_{E,C}^{\Pi^{cov}}(\kappa)$ denote the security experiment. The main property of covert QKE, denoted by Π^{cov} , can be expressed as a game played by the adversary E against a challenger C who decides on a random bit b and if $b = 0$, C runs Π^{cov} , otherwise (if $b = 1$), C does not run Π^{cov} . Finally, E guesses a random bit b' and wins the game if $b = b'$. The experiment outputs 1 if E succeeds, and 0 otherwise. The winning advantage of E is given by $\text{Adv}_E^{\Pi^{cov}}(\kappa) = |\Pr[b = b'] - 1/2|$ and we want that $\text{Adv}_E^{\Pi^{cov}}(\kappa) \leq \text{negl}(\kappa)$.

Definition 2. Let $G : \{0, 1\}^s \rightarrow \{0, 1\}^{g(s)}$ be a (τ, ϵ) -PRNG secure against all efficient distinguishers D running in time at most τ with success probability at most ϵ , where $\forall s : g(s) > s$. A QKE protocol Π_G^{cov} is considered to be covert if the following holds for any efficient adversary E :

- Π_G^{cov} is a secure QKE protocol.
- The probability that E guesses the bit b correctly ($b' = b$), i.e., E manages to distinguish between Alice and Bob running Π_G^{cov} or not, is no more than $\frac{1}{2}$ plus a negligible function in the security parameter κ , i.e.,

$$\Pr[\text{CovertQKE}_{E,C}^{\Pi^{cov}}(\kappa) = 1] \leq \frac{1}{2} + \text{negl}(\kappa)$$

Theorem 1. (Sourced from [1]) The secret key obtained from the covert QKE protocol Π_G^{cov} is informational-theoretically secure and the covertness of Π_G^{cov} is as secure as the underlying PRNG.

4.2 Deniable Covert Quantum Key Exchange (DC-QKE)

We are now in a position to describe DC-QKE, a simple construction shown in Protocol 4, which preserves unconditional security for the final secret key, while its deniability is as secure as the underlying PRNG used in $\Pi_{r,G}^{cov}$. In terms of the Security Experiment 3.4, $\Pi_{r,G}^{cov}$ is run to establish a real key k , while non-covert QKE $\Pi_{r'}$ is used to produce a fake key k' aimed at achieving deniability, where r and r' are the respective vectors of real and fake private inputs.

Operationally, consider a setting wherein the parties suspect in advance that they might be coerced into revealing their private coins for a given run: their joint strategy consists of running both components in Protocol 4 and claiming to have employed $\Pi_{r'}$ to establish the fake key k' using the fake private randomness r' (e.g. raw key bits in BB84) and provide these as input to the adversary upon termination of a session. Thus, for Eve to be able to produce a proof showing that the revealed values are fake, she would have to break the security of covert QKE to detect the presence of $\Pi_{r,G}^{cov}$, as shown in Theorem 2. Moreover, note that covert communication can be used for dynamically agreeing on a joint strategy for denial, further highlighting its relevance for deniability.

Protocol 4 DC-QKE for an n -bit key

- 1: **RandGen:** Let $\mathbf{r} = (r_1, \dots, r_\ell)$ be the vector of private random inputs, where $r_i \leftarrow_{\$} \{0, 1\}^{|r_i|}$.
 - 2: **KeyGen:** Run $\Pi_{\mathbf{r}, G}^{cov}$ to establish a random secret key $k \in \{0, 1\}^n$.
- Non-covert faking component $\mathcal{F}_{A,B}$:
- 1: **FakeRandGen:** Let $\mathbf{r}' = (r'_1, \dots, r'_\ell)$ be the vector of fake private random inputs, where $r'_i \leftarrow_{\$} \{0, 1\}^{|r'_i|}$.
 - 2: **FakeKeyGen:** Run $\Pi_{\mathbf{r}'}$ to establish a separate fake key $k' \in \{0, 1\}^n$.
-

Remark 2. The original analysis in [3] describes an attack based solely on revealing fake raw key bits that may be inconsistent with the adversary's observations. An advantage of DC-QKE in this regard is that Alice's strategy for achieving coercion-deniability consists of revealing all the secret values of the non-covert QKE $\Pi_{\mathbf{r}'}$ honestly. This allows her to cover the full range of private randomness that could be considered in different variants of deniability as discussed in Remark 1. A potential drawback is the extra cost induced by $\mathcal{F}_{A,B}$, which could, in principle, be mitigated using a less interactive solution such as QKE via UE.

Remark 3. If the classical channel is authenticated by an information-theoretically secure algorithm, the minimal entropy overhead in terms of pre-shared key (logarithmic in the input size) for Π can be generated by $\Pi_{\mathbf{r}}^{cov}$.

Example 1. In the case of encryption, A can send $c = m \oplus k$ over a covert channel to B , while for denying to m' , she can send $c' = m' \oplus k'$ over a non-covert channel. Alternatively, she can transmit a single ciphertext over a non-covert channel such that it can be opened to two different messages. To do so, given $c = m \oplus k$, Alice computes $k' = m' \oplus c = m' \oplus m \oplus k$, and she can then either encode k' as a codeword, as described in Section 2.1, and run $\Pi_{\mathbf{r}'}$ via uncloneable encryption, thus allowing her to reveal the entire transcript to Eve honestly, or she can agree with Bob on a suitable privacy amplification (PA) function (with PA being many-to-one) as part of their denying program in order to obtain k' .

Theorem 2. *If $\Pi_{\mathbf{r}, G}^{cov}$ is a covert QKE protocol, then DC-QKE given in Protocol 4 is a coercion-deniable QKE protocol that satisfies Definition 1.*

Proof. The main idea consists of showing that breaking the deniability property of DC-QKE amounts to breaking the security of covert QKE, such that coercion-deniability follows from the contrapositive of this implication, i.e., if there exists no efficient algorithm for compromising the security of covert QKE, then there exists no efficient algorithm for breaking the deniability of DC-QKE. We formalize this via a reduction, sketched as follows. Let $w' = \text{View}_{Fake}(k', \mathbf{t}_E, \mathbf{r}')$ and $w = \text{View}_{Real}(k, \mathbf{t}_E, \mathbf{r})$ denote the two views. Flip a coin b for an attempt at denial: if $b = 0$, then $\mathbf{t}_E = (\mathbf{t}', \emptyset)$, else ($b = 1$), $\mathbf{t}_E = (\mathbf{t}', \mathbf{t}^{cov})$, where \mathbf{t}^{cov} and \mathbf{t}' denote the transcripts of covert and non-covert exchanges from $\Pi_{\mathbf{r}, G}^{cov}$ and $\Pi_{\mathbf{r}'}$. Now if DC-QKE is constructed from $\Pi_{\mathbf{r}, G}^{cov}$, then given an efficient adversary E that can distinguish w from w' with probability p_1 , we can use E to construct an efficient distinguisher D to break the security of

covert QKE with probability p_2 such that $p_1 \leq p_2$. Indeed, given an instance of a DC-QKE security game, we construct a distinguisher D that uses E on input w and w' , with the goal to win the game described in the Security Experiment 3.4. The distinguisher D would simply run E (with negligible overhead) and observe whether E succeeds at distinguishing w from w' . Since the only element that is not sampled uniformly at random is in \mathbf{t}^{cov} containing exchanges from the covert channel, which relies on a PRNG, the only way E can distinguish w from w' is if she can distinguish $(\mathbf{t}', \mathbf{t}^{cov})$ from (\mathbf{t}', \emptyset) . If E succeeds, then D guesses that a covert QKE session has taken place, thereby winning the Security Experiment 4.1 for covert QKE. \square

5 Deniability via Entanglement Distillation

Here we consider the possibility of achieving information-theoretic deniability via entanglement distillation (ED). In its most general form, ED allows two parties to distill maximally entangled pure states (*ebits*) from an arbitrary sequence of entangled states at some positive rate using local operations and classical communication (LOCC), i.e. to move from $|\Phi_\theta\rangle_{AB} \equiv \cos(\theta)|00\rangle_{AB} + \sin(\theta)|11\rangle_{AB}$ to $|\Phi^+\rangle_{AB} = \frac{1}{\sqrt{2}}(|00\rangle_{AB} + |11\rangle_{AB})$, where $0 < \theta < \pi/2$.

In the noiseless model, n independent identically distributed (i.i.d.) copies of the same partially entangled state ρ can be converted into $\approx nH(\rho)$ Bell pairs in the limit $n \rightarrow \infty$, i.e., from $\rho_{AB}^{\otimes n}$ to $|\Phi^+\rangle_{AB}^{\otimes nH(\rho)}$, where $H(\rho) = -\text{Tr}(\rho \ln \rho)$ denotes the von Neumann entropy of entanglement. If the parties start out with pure states, local operations alone will suffice for distillation [4,7], otherwise the same task can be achieved via forward classical communication (one-way LOCC), as shown by the Devetak-Winter theorem [14], to distill ebits from many copies of some bipartite entangled state. See also the early work of Bennett et al. [8] on mixed state ED. Buscemi and Datta [9] relax the i.i.d. assumption and provide a general formula for the optimal rate at which ebits can be distilled from a noisy and arbitrary source of entanglement via one-way and two-way LOCC.

Intuitively, the eavesdropping attack described in [3] and further detailed in Section 3.1, is enabled by the presence of noise in the channel as well as the fact that Bob cannot distinguish states sent by Alice from those prepared by Eve. As a result, attempting to deny to a different bit value encoded in a given quantum state - without knowing if this is a decoy state prepared by Eve - allows the adversary to detect such an attempt with non-negligible probability.

In terms of deniability, the intuition behind this idea is that while Alice and Bob may not be able to know which states have been prepared by Eve, they can instead remove her “check” decoy states from their set of shared entangled pairs by decoupling her system from theirs. Once they are in possession of maximally entangled states, they will have effectively factored out Eve’s state such that the global system is given by the pure tensor product space $|\Psi^+\rangle_{AB} \otimes |\phi\rangle_E$. Thus the pure bipartite joint system between Alice and Bob cannot be correlated with any system under Eve’s control, thereby foiling her cross-checking strategy. The singlet states can then be used to perform QKE via quantum teleportation [6].

5.1 Deniable QKE via Entanglement Distillation and Teleportation

We now argue why performing randomness distillation at the quantum level, thus requiring quantum computation, plays an important role w.r.t. deniability. The subtleties alluded to in [3] arise from the fact that randomness distillation is performed in the classical post-processing step. This allows Eve to leverage her tampering in that she can verify the parties' claims against her decoy states. However, this attack can be countered by removing Eve's knowledge before the classical exchanges begin. Most security proofs of QKE [22,28,23] are based on a reduction to an entanglement-based variant, such that the fidelity of Alice and Bob's final state with $|\Psi^+\rangle^{\otimes m}$ is shown to be exponentially close to 1. Moreover, secret key distillation techniques involving ED and quantum teleportation [7,14] can be used to faithfully transfer qubits from A to B by consuming ebits. To illustrate the relevance of distillation for deniability in QKE, consider the generalized template shown in Protocol 5, based on these well-known techniques.

Protocol 5 Template for deniable QKE via entanglement distillation and teleportation

- 1: A and B share n noisy entangled pairs (assume i.i.d. states for simplicity).
 - 2: They perform entanglement distillation to convert them into a state ρ such that $F(|\Psi^+\rangle^{\otimes m}, \rho)$ is arbitrarily close to 1 where $m < n$.
 - 3: Perform verification to make sure they share m maximally entangled states $|\Psi^+\rangle^{\otimes m}$, and abort otherwise.
 - 4: A prepares m qubits (e.g. BB84 states) and performs quantum teleportation to send them to B at the cost of consuming m ebits and exchanging $2m$ classical bits.
 - 5: A and B proceed with standard classical distillation techniques to agree on a key based on their measurements.
-

By performing ED, Alice and Bob make sure that the resulting state cannot be correlated with anything else due to the monogamy of entanglement (see e.g. [21,30]), thus factoring out Eve's system. The parties can open their records for steps (2) and (3) honestly, and open to arbitrary classical inputs for steps (3), (4) and (5): deniability follows from decoupling Eve's system, meaning that she is faced with a reduced density matrix on a pure bipartite maximally entangled state, i.e., a maximally mixed state $\rho_E = \mathbb{I}/2$, thus obtaining key equivocation.

In terms of the hierarchy of entanglement-based constructions mentioned in [3], this approach mainly constitutes a generalization of such schemes. It should therefore be viewed more as a step towards a theoretical characterization of entanglement-based schemes for achieving information-theoretic deniability. Due to lack of space, we omit a discussion of how techniques from device-independent cryptography can deal with maliciously prepared initial states.

Going beyond QKE, note that quantum teleportation allows the transfer of an *unknown* quantum state, meaning that even the sender would be oblivious as to what state is sent. Moreover, ebits can enable uniquely quantum tasks such as *traceless exchange* in the context of quantum anonymous transmission [12], to achieve *incoercible* protocols that allow parties to deny to any random input.

6 Open Questions and Directions for Future Research

Studying the deniability of public-key authenticated QKE both in our model and in the simulation paradigm, and the existence of an equivalence relation between our indistinguishability-based definition and a simulation-based one would be a natural continuation of this work. Other lines of inquiry include forward deniability, deniable QKE in conjunction with forward secrecy, deniability using covert communication in stronger adversarial models, a further analysis of the relation between the impossibility of unconditional quantum bit commitment and deniability mentioned in [3], and deniable QKE via uncloneable encryption. Finally, gaining a better understanding of entanglement distillation w.r.t. potential pitfalls in various adversarial settings and proposing concrete deniable protocols for QKE and other tasks beyond key exchange represent further research avenues.

Acknowledgments

We thank Mark M. Wilde and Ignatius William Primaatmaja for their comments. This work was supported by a grant (Q-CoDe) from the Luxembourg FNR.

References

1. Arrazola, J.M., Scarani, V.: Covert quantum communication. *Physical review letters* **117**(25), 250503 (2016)
2. Bash, B.A., Goeckel, D., Towsley, D., Guha, S.: Hiding information in noise: Fundamental limits of covert wireless communication. *IEEE Communications Magazine* **53**(12), 26–31 (2015)
3. Beaver, D.: On deniability in quantum key exchange. In: Knudsen, L.R. (ed.) *Advances in Cryptology — EUROCRYPT 2002*. pp. 352–367. Springer Berlin Heidelberg, Berlin, Heidelberg (2002)
4. Bennett, C.H., Bernstein, H.J., Popescu, S., Schumacher, B.: Concentrating partial entanglement by local operations. *Physical Review A* **53**(4), 2046 (1996)
5. Bennett, C.H., Brassard, G.: Quantum cryptography: public key distribution and coin tossing int. In: *Conf. on Computers, Systems and Signal Processing (Bangalore, India, Dec. 1984)*. pp. 175–9 (1984)
6. Bennett, C.H., Brassard, G., Crépeau, C., Jozsa, R., Peres, A., Wootters, W.K.: Teleporting an unknown quantum state via dual classical and einstein-podolsky-rosen channels. *Physical review letters* **70**(13), 1895 (1993)
7. Bennett, C.H., Brassard, G., Popescu, S., Schumacher, B., Smolin, J.A., Wootters, W.K.: Purification of noisy entanglement and faithful teleportation via noisy channels. *Physical review letters* **76**(5), 722 (1996)
8. Bennett, C.H., DiVincenzo, D.P., Smolin, J.A., Wootters, W.K.: Mixed-state entanglement and quantum error correction. *Physical Review A* **54**(5), 3824 (1996)
9. Buscemi, F., Datta, N.: Distilling entanglement from arbitrary resources. *Journal of Mathematical Physics* **51**(10), 102201 (2010)
10. Calderbank, A.R., Shor, P.W.: Good quantum error-correcting codes exist. *Physical Review A* **54**(2), 1098 (1996)
11. Canetti, R., Dwork, C., Naor, M., Ostrovsky, R.: Deniable encryption. In: *Annual International Cryptology Conference*. pp. 90–104. Springer (1997)

12. Christandl, M., Wehner, S.: Quantum anonymous transmissions. In: International Conference on the Theory and Application of Cryptology and Information Security. pp. 217–235. Springer (2005)
13. Cremers, C., Feltz, M.: One-round strongly secure key exchange with perfect forward secrecy and deniability. Tech. rep., ETH Zurich (2011)
14. Devetak, I., Winter, A.: Distillation of secret key and entanglement from quantum states. *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences* **461**(2053), 207–235 (2005)
15. Di Raimondo, M., Gennaro, R.: New approaches for deniable authentication. *Journal of cryptology* **22**(4), 572–615 (2009)
16. Di Raimondo, M., Gennaro, R., Krawczyk, H.: Deniable authentication and key exchange. In: *Proceedings of the 13th ACM conference on Computer and communications security*. pp. 400–409. ACM (2006)
17. Dodis, Y., Katz, J., Smith, A., Walfish, S.: Composability and on-line deniability of authentication. In: *Theory of Cryptography Conference*. pp. 146–162. Springer (2009)
18. Dwork, C., Naor, M., Sahai, A.: Concurrent zero-knowledge. In: *Proceedings of the 30th Annual ACM Symposium on Theory of Computing*. pp. 409–418. STOC '98, ACM, New York, NY, USA (1998)
19. Gottesman, D.: Uncloneable encryption. *Quantum Info. Comput.* **3**(6), 581–602 (Nov 2003)
20. Ioannou, L.M., Mosca, M.: A new spin on quantum cryptography: Avoiding trapdoors and embracing public keys. In: *International Workshop on Post-Quantum Cryptography*. pp. 255–274. Springer (2011)
21. Koashi, M., Winter, A.: Monogamy of quantum entanglement and other correlations. *Physical Review A* **69**(2), 022309 (2004)
22. Lo, H.K., Chau, H.F.: Unconditional security of quantum key distribution over arbitrarily long distances. *science* **283**(5410), 2050–2056 (1999)
23. Mayers, D.: Unconditional security in quantum cryptography. *Journal of the ACM (JACM)* **48**(3), 351–406 (2001)
24. Mosca, M., Stebila, D., Ustaoglu, B.: Quantum key distribution in the classical authenticated key exchange framework. In: *International Workshop on Post-Quantum Cryptography*. pp. 136–154. Springer (2013)
25. Nielsen, M.A., Chuang, I.: *Quantum computation and quantum information* (2002)
26. Pass, R.: On deniability in the common reference string and random oracle model. In: *Annual International Cryptology Conference*. pp. 316–337. Springer (2003)
27. Sheikholeslami, A., Bash, B.A., Towsley, D., Goeckel, D., Guha, S.: Covert communication over classical-quantum channels. In: *Information Theory (ISIT), 2016 IEEE International Symposium on*. pp. 2064–2068. IEEE (2016)
28. Shor, P.W., Preskill, J.: Simple proof of security of the bb84 quantum key distribution protocol. *Physical review letters* **85**(2), 441 (2000)
29. Steane, A.: Multiple-particle interference and quantum error correction. *Proc. R. Soc. Lond. A* **452**(1954), 2551–2577 (1996)
30. Streltsov, A., Adesso, G., Piani, M., Bruß, D.: Are general quantum correlations monogamous? *Physical review letters* **109**(5), 050503 (2012)
31. Unger, N., Goldberg, I.: Deniable key exchanges for secure messaging. In: *Proceedings of the 22nd acm sigsac conference on computer and communications security*. pp. 1211–1223. ACM (2015)
32. Wilde, M.M.: *Quantum information theory*. Cambridge University Press (2013)
33. Wootters, W.K., Zurek, W.H.: A single quantum cannot be cloned. *Nature* **299**(5886), 802–803 (1982)