

Computational Security of Quantum Encryption

Gorjan Alagic¹, Anne Broadbent², Bill Fefferman³, Tommaso Gagliardoni⁴,
Christian Schaffner⁵, and Michael St. Jules⁶

¹ Department of Mathematical Sciences, University of Copenhagen
`galagic@gmail.com`

² Department of Mathematics and Statistics, University of Ottawa
`abroadbe@uottawa.ca`

³ QuICS, University of Maryland/NIST
`wjf@umd.edu`

⁴ Cryptoplexity, TU Darmstadt, Germany
`tommaso@gagliardoni.net`

⁵ QuSoft, University of Amsterdam and CWI, The Netherlands
`c.schaffner@uva.nl`

⁶ Department of Mathematics and Statistics, University of Ottawa
`mstju032@uottawa.ca`

Abstract. Quantum-mechanical devices have the potential to transform cryptography. Most research in this area has focused either on the information-theoretic advantages of quantum protocols or on the security of classical cryptographic schemes against quantum attacks. In this work, we initiate the study of another relevant topic: the encryption of quantum data in the computational setting.

In this direction, we establish quantum versions of several fundamental classical results. First, we develop natural definitions for private-key and public-key encryption schemes for quantum data. We then define notions of semantic security and indistinguishability, and, in analogy with the classical work of Goldwasser and Micali, show that these notions are equivalent. Finally, we construct secure quantum encryption schemes from basic primitives. In particular, we show that quantum-secure one-way functions imply IND-CCA1-secure symmetric-key quantum encryption, and that quantum-secure trapdoor one-way permutations imply semantically-secure public-key quantum encryption.

Table of Contents

1	Introduction.....	1
1.1	Summary of Contributions and Techniques	1
1.2	Related Work.....	4
2	Preliminaries	5
2.1	Classical States, Maps, and the One-Time Pad	5
2.2	Quantum States, Maps, and the One-Time Pad	6
2.3	Efficient Classical and Quantum Computations.....	7
2.4	Oracles	9
3	Quantum Encryption and Indistinguishability	9
3.1	Quantum Encryption Schemes	10
3.2	Indistinguishability of Encryptions.....	11
4	Quantum Semantic Security	12
4.1	Difficulties in the Quantum Setting	12
4.2	Definition of Semantic Security.....	13
4.3	Semantic Security is Equivalent to Indistinguishability	14
5	Quantum Encryption Schemes	16
5.1	Quantum Symmetric-Key Encryption from One-Way Functions ..	16
5.2	Quantum Public-Key Encryption from Trapdoor Permutations ..	18
6	Conclusion	22
6.1	Extensions and Future Work.....	22
6.2	Acknowledgements	23
A	Alternative Definitions of Quantum Security.....	26
A.1	SEM2.....	27
A.2	SEM3.....	28
A.3	IND'.....	29

1 Introduction

Quantum mechanics changes our view of information processing: the ability to access, operate and transmit data according to the laws of quantum physics opens the doors to a vast realm of possible applications. Cryptography is one of the areas that is most seriously impacted by the potential of quantum information processing, since the security of most cryptographic primitives in use today relies on the hardness of computational problems that are easily broken by adversaries having access to a quantum computer [Sho94].

While the impact of quantum computers on cryptanalysis is tremendous, quantum mechanics itself predicts physical phenomena that can be exploited in order to achieve new levels of security. These advantages were already mentioned in the late 1970’s in pioneering work of Wiesner [Wie83], and have led to the very successful theory of quantum key distribution (QKD) [BB84], which has already seen real-world applications [ABB⁺14]. QKD achieves information-theoretically secure key expansion, and has the advantage of relatively simple hardware requirements (notwithstanding a long history of successful attacks to QKD at the implementation level [ABB⁺14]).

The cryptographic possibilities of quantum information go well beyond QKD. Indeed, quantum copy-protection [Aar09], quantum money [Wie83,AC12,MS10] and revocable time-release encryption [Unr14] are just some examples where properties unique to quantum data enable new cryptographic constructions (see [BS16] for a survey). Thanks in part to these tremendous cryptographic opportunities, we envisage an increasing need for an information infrastructure that enables quantum information. Such an infrastructure will be required to support:

- **Quantum functionality:** honest parties can store, exchange, and compute on quantum data;
- **Quantum security:** quantum functionality is protected against quantum adversaries.

The current state-of-the-art is lacking even the most basic cryptographic concepts in the context of quantum functionality and quantum adversaries. In particular, the study of encryption of quantum data (which is arguably one of the most fundamental building blocks) has so far been almost exclusively limited to the quantum one-time pad [AMTdW00] and other aspects of the information-theoretic setting [Des09,DD10] (one notable exception being [BJ15]). The achievability of other basic primitives such as public-key encryption has not been thoroughly investigated for the case of fully quantum cryptography. This situation leaves many open questions about what can be achieved in the quantum world.

1.1 Summary of Contributions and Techniques

In this work, we establish quantum versions of several fundamental classical (*i.e.* “non-quantum”) results in the setting of computational security. Following

Broadbent and Jeffrey [BJ15], we consider private-key and public-key encryption schemes for quantum data. In these schemes, the key is a classical bitstring⁷, but both the plaintext and the ciphertext are quantum states. Key generation, encryption, and decryption are implemented by polynomial-time quantum algorithms. Such schemes admit an appropriate definition of indistinguishability security, following the classical approach [BJ15]: the quantum adversary is given access to an encryption oracle, and must output a challenge plaintext; given either the corresponding ciphertext or the encryption of $|0\rangle\langle 0|$ (each with probability $1/2$), the adversary must decide which was the case.

Our main contributions are the following. First, we give several natural formulations of semantic security for quantum encryption schemes, and show that all of them are equivalent to indistinguishability. This cements the intuition that possession of the ciphertext should not help the adversary in computing anything about the plaintext. Second, we give two constructions of encryption schemes with semantic security: a private-key scheme, and a public-key scheme. The private-key scheme satisfies a stronger notion of security: indistinguishability against chosen ciphertext attacks (IND-CCA1). A more detailed summary of these contributions follows.

1.1.1 Semantic Security vs. Indistinguishability Semantic security formalizes the notion of security of an encryption scheme under computational assumptions. Originally introduced by Goldwasser and Micali [GM84], this definition posits a game: an adversary is given the encryption of a message x and some side information $h(x)$, and is challenged to output the value of an objective function f evaluated at x . An encryption scheme is deemed secure if every adversary can be closely approximated by a *simulator* who is given only $h(x)$; crucially, the simulator must work for every possible choice (h, f) of side information and objective function. This models the intuitive notion that having access to a ciphertext gives the adversary essentially no advantage in computing functions related to the plaintext.

While semantic security corresponds to a notion of security that is intuitively strong, it is cumbersome to use in terms of security proofs. In order to address this problem, Goldwasser and Micali [GM84] showed the equivalence of semantic security with another cryptographic notion, called *indistinguishability*. The intuitive description of indistinguishability is also in terms of a game, this time with a *single* adversary. The adversary prepares a pair of plaintexts x_0 and x_1 and submits them to a challenger, who chooses a uniformly random bit b and returns the encryption of x_b . The adversary then performs a computation and outputs a bit v ; the adversary wins the game if $v = b$ and loses otherwise. An encryption scheme is deemed secure if no adversary wins the game with probability significantly larger than $1/2$. This definition models the intuitive notion that the ciphertexts are indistinguishable: whatever the adversary does with one ciphertext, the outcome is essentially the same if run on the other ciphertext.

⁷ While quantum keys might be of interest, they are not necessary for constructing secure schemes [BJ15].

In [Section 4](#), we define semantic security for the encryption of *quantum* data—thus establishing a parallel with the notions and results of encryptions as laid out by Goldwasser and Micali. When attempting to transfer the definition of semantic security to the quantum world, the main question one encounters is to determine the quantum equivalents of $h(x)$ and $f(x)$ as described above (because of the no-cloning theorem [\[WZ82\]](#), we cannot postulate a polynomial-time experiment that simultaneously involves some quantum plaintext *and* a function of the plaintext—see [Section 4.2](#) for further discussions related to this issue). We propose a number of alternative definitions in order to deal with this situation ([Definition 8](#), [Definition 22](#), and [Definition 25](#).) Perhaps the most surprising is our definition of SEM ([Definition 8](#)), which does away completely with the need to explicitly define analogues of the functions h and f , instead relying on a *message generator* that outputs three registers, consisting of the “plaintext”, “side information” and “target output” (there is no further structure imposed on the contents of these registers). Intuitively, we think of the adversary’s goal being to output the value contained in the “target output” register. Formally, however, [Definition 8](#) shows that the role of the “target output” register is actually to help the distinguisher: semantic security corresponding to the situation where no distinguisher has a non-negligible advantage in telling apart the real scenario (involving the adversary) and the ideal scenario (involving the simulator), *even given access to the “target output” system*. Our main result in this direction (see [Section 4.3](#)) is the equivalence between semantic security and indistinguishability for quantum encryption schemes:

Theorem 1. *A quantum encryption scheme is semantically secure if and only if it has indistinguishable encryptions.*

What is more, because our definitions and proofs hold when restricted to the classical case (and in fact can be shown as generalizations of the standard classical definitions), our contribution sheds new light on semantic security: to the best of our knowledge, this is the first time that semantic security has been defined *without* the need to explicitly refer to functions h and f .

1.1.2 Quantum Encryption Schemes In [Section 5](#), we give two constructions of quantum encryption schemes that achieve semantic security (and thus also indistinguishability, by [Theorem 1](#).) Our constructions make use of two basic primitives. The first is a *quantum-secure one-way function* (qOWF). This is a family of deterministic functions which are efficiently computable in classical polynomial time, but which are impossible to invert even in quantum polynomial time. It is believed that such functions can be constructed from certain algebraic problems [\[MRV07, KK07\]](#). The existence of qOWFs implies the existence of *quantum-secure pseudorandom functions* (qPRFs) [\[Zha12\]](#). We show that a qPRF can, in turn, be used to securely encrypt quantum data with classical private keys. More precisely, we have the following:

Theorem 2. *If quantum-secure one-way functions exist, then so do IND-CCA1-secure private-key quantum encryption schemes.*

The second basic primitive we consider is a *quantum-secure one-way permutation with trapdoors* (qTOWP). In analogy with the classical case, a qTOWP is a qOWF with an additional property: each function in the family is a permutation whose efficient inversion is possible if one possesses a secret string (the trapdoor). While our results appear to be the first to consider applications to quantum data, the notion of quantum security for trapdoor permutations is of obvious relevance in the security of classical cryptosystems against quantum attacks. Some promising candidate qTOWPs from lattice problems are known [PW08,GPV08]. We show that such functions can be used to give secure public-key encryption schemes for quantum data, again using only classical keys.

Theorem 3. *If quantum-secure trapdoor one-way permutations exist, then so do semantically secure public-key quantum encryption schemes.*

We remark that Theorem 2 and Theorem 3 are analogues of standard results in the classical literature [Gol04a].

1.2 Related Work

Prior work has considered the computational security of quantum methods to encrypt classical data [OTU00,Kos07,XY12]. Information-theoretic security for the encryption of quantum states has been considered in the context of the one-time pad [AMTdW00,BR03,HLSW04,Leu02], as well as entropic security [Des09,DD10]. Computational indistinguishability notions for encryption in a quantum world were proposed in independent and concurrent work [BJ15,GHS15]. While [BJ15] considers the encryption of quantum data (and proposes the first constructions based on hybrid classical-quantum encryption), [GHS15] considers the security of *classical* schemes which can be accessed in a quantum way by the adversary.

The results of [GHS15] are part of a line of research of “*post-quantum*” cryptography, which investigates the security of classical schemes against quantum adversaries, with the goal of finding “quantum-safe” schemes. This includes the study of encryption and signature schemes secure against attacks by quantum algorithms [BBD09], and also the study of superposition attacks against quantum oracles [BDF⁺11,Zha12,Unr15]. Still in the model of superposition attacks, [BZ13] studies quantum indistinguishability under chosen plaintext and chosen ciphertext attacks. This definition was improved in [GHS15] to allow for a quantum challenge phase. The latter paper also initiates the study of quantum semantic security of classical schemes and gives the first classical construction of a quantumly secure encryption scheme from a family of quantum-secure pseudo-random permutations. Another quantum indistinguishability notion in the same spirit has been suggested (but not further analyzed) in [Vel13, Def. 5.3].

Several previous works have considered how classical security proofs change in the setting of quantum attacks (see, e.g., [Unr10,FKS⁺13,Son14].) Our results can be viewed as part of this line of work; one distinguishing feature is that we are able to extend classical security proofs to the setting of quantum functionality secure against quantum adversaries. This setting has seen increasing interest in the

past decade, with progress being made on several topics: multi-party quantum computation [BOCG⁺06], secure function evaluation [DNS10,DNS12], one-time programs [BGS13], and delegated quantum computation [BFK09,Bro15].

Outline. The remainder of the paper is structured as follows. In Section 2, we set down basic notation and recall a few standard facts regarding classical and quantum computation. In Section 3, we define symmetric-key and public-key encryption for quantum states (henceforth “quantum encryption schemes”), as well as a notion of indistinguishability (including IND-CPA and IND-CCA1) for such schemes. Section 4 defines semantic security for quantum encryption schemes, and shows equivalence with indistinguishability. Section 5 gives our two constructions for quantum encryption schemes. Finally, we close with some discussion of future work in Section 6.

2 Preliminaries

We introduce some basic notation for classical (Section 2.1) and quantum (Section 2.2) information processing and information-theoretic encryption. Section 2.3 concerns basic issues in efficient algorithms and Section 2.4 discusses the use of oracles.

2.1 Classical States, Maps, and the One-Time Pad

Let \mathbb{N} be the set of positive integers. For $n \in \mathbb{N}$, we set $[n] = \{1, \dots, n\}$. Define $\{0, 1\}^* := \cup_n \{0, 1\}^n$. An element $x \in \{0, 1\}^*$ is called a bitstring, and $|x|$ denotes its length, *i.e.*, its number of bits. We reserve the notation 0^n (resp., 1^n) to denote the n -bit string with all zeroes (resp., all ones).

For a finite set X , the notation $x \stackrel{\$}{\leftarrow} X$ indicates that x is selected uniformly at random from X . For a probability distribution S , the notation $x \leftarrow S$ indicates that x is sampled according to S . Given finite sets X and Y , the set of all functions from Y to X is denoted X^Y (or sometimes $\{X \rightarrow Y\}$). We will usually consider functions f acting on binary strings, that is, of the form $f : \{0, 1\}^n \rightarrow \{0, 1\}^m$, for some positive integers n and m . We will also consider function families $f : \{0, 1\}^* \rightarrow \{0, 1\}^*$ defined on bitstrings of arbitrary size. One can construct such a family simply by choosing one function with input size n , for each n . We will sometimes abuse notation by stating that $f : \{0, 1\}^n \rightarrow \{0, 1\}^m$ defines a function family; in that case, it is implicit that n is a parameter that indexes the input size and m is some function of n (usually a polynomial) that indexes the output size. Given a bitstring y and a function family f , the preimage of f under y is defined by $f^{-1}(y) := \{x \in \{0, 1\}^* : f(x) = y\}$.

We will often write $\text{negl}(\cdot)$ to denote a function from \mathbb{N} to \mathbb{N} which is “negligible” in the sense that it grows at an inverse-superpolynomial rate. More precisely, $\text{negl}(n) < 1/p(n)$ for every polynomial $p : \mathbb{N} \rightarrow \mathbb{N}$ and all sufficiently large n . A typical use of negligible functions is to indicate that the probability of success

of some algorithm is too small to be amplified to a constant by a feasible (*i.e.*, polynomial) number of repetitions.

Given two bitstrings x and y of equal length, we denote their bitwise XOR by $x \oplus y$. Recall that the *classical one-time pad* encrypts a plaintext $x \in \{0, 1\}^n$ by XORing it with a uniformly random string (the key) $r \xleftarrow{\$} \{0, 1\}^n$. Decryption is performed by repeating the operation, *i.e.*, by XORing the key with the ciphertext. Since the uniform distribution on $\{0, 1\}^n$ is invariant under XOR by x , the ciphertext is uniformly random to parties having no knowledge about r [Sha49]. A significant drawback of the one-time pad is the key length. In order to reduce the key length, one may generate r pseudorandomly; this key-length reduction requires making computational assumptions about the adversary.

2.2 Quantum States, Maps, and the One-Time Pad

Given an n -bit string x , the corresponding quantum-computational n -qubit basis state is denoted $|x\rangle$. The 2^n -dimensional Hilbert space spanned by n -qubit basis states will be denoted

$$\mathcal{H}_n := \text{span} \{|x\rangle : x \in \{0, 1\}^n\} .$$

We denote by $\mathfrak{D}(\mathcal{H}_n)$ the set of density operators (*i.e.*, valid quantum states) on \mathcal{H}_n . These are linear operators on $\mathfrak{D}(\mathcal{H}_n)$ which are positive-semidefinite and have trace equal to 1. When considering different physical subsystems, we will denote them with uppercase Latin letters; when a Hilbert space corresponds to a subsystem, we will place the subsystem label in the subscript. For instance, if $F \cup G \cup H = [n]$ then $\mathcal{H}_n = \mathcal{H}_F \otimes \mathcal{H}_G \otimes \mathcal{H}_H$. Sometimes we will write explicitly the subsystems a state belongs to as subscripts; this will be useful when considering, *e.g.*, the reduced state on some of the subspaces. For example, we will sometimes express the statement $\rho \in \mathfrak{D}(\mathcal{H}_F \otimes \mathcal{H}_G \otimes \mathcal{H}_H)$ simply by calling the state ρ_{FGH} ; in that case, the state obtained by tracing out the subsystem H will be denoted ρ_{FG} .

Given $\rho, \sigma \in \mathfrak{D}(\mathcal{H})$, the trace distance between ρ and σ is given by half the trace norm $\|\rho - \sigma\|_1$ of their difference. When ρ and σ are classical probability distributions, the trace distance reduces to the total variation distance. Physically realizable maps from a state space $\mathfrak{D}(\mathcal{H})$ to another state space $\mathfrak{D}(\mathcal{H}')$ are called *admissible*—these are the completely positive trace-preserving (CPTP) maps. For the purpose of distinguishability via input/output operations, the appropriate norm for CPTP maps is the diamond norm, denoted $\|\cdot\|_\diamond$. The set of admissible maps coincides with the set of all maps realizable by composing (i.) addition of ancillas, (ii.) unitary evolutions, (iii.) measurements in the computational basis, and (iv.) tracing out subspaces. We remark that unitaries $U \in U(\mathcal{H}_n)$ act on $\mathfrak{D}(\mathcal{H}_n)$ by conjugation: $\rho \mapsto U\rho U^\dagger$. The identity operator $\mathbb{1}_n \in U(\mathcal{H}_n)$ is thus both a valid map, and (when normalized by 2^{-n}) a valid state in $\mathfrak{D}(\mathcal{H}_n)$ —corresponding to the classical uniform distribution.

Recall the single-qubit Pauli operators defined as:

$$I = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \quad X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad Y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \quad Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}.$$

The Pauli operators are Hermitian and unitary quantum gates, i.e. $P^\dagger = P$ and $P^\dagger P = P P^\dagger = P^2 = I$ for all $P \in \{I, X, Y, Z\}$. It is easy to check that applying a uniformly random Pauli operator to any single-qubit density operator results in the maximally mixed state:

$$\frac{1}{4}(\rho + X\rho X + Y\rho Y + Z\rho Z) = \frac{\mathbb{1}_1}{2} \quad \text{for all } \rho \in \mathfrak{D}(\mathcal{H}_1).$$

Since the Pauli operators are self-adjoint, we may implement the above map by choosing two bits s and t uniformly at random and then applying

$$\rho \mapsto X^s Z^t \rho Z^t X^s.$$

To observers with no knowledge of s and t , the resulting state is information-theoretically indistinguishable from $\mathbb{1}_1/2$. Of course, if we know s and t , we can invert the above map and recover ρ completely.

The above map can be straightforwardly extended to the n -qubit case in order to obtain an elementary *quantum encryption scheme* called the *quantum one-time pad*. We first set $X_j = \mathbb{1}^{\otimes j-1} \otimes X \otimes \mathbb{1}^{\otimes n-j}$ and likewise for Y_j and Z_j . We define the n -qubit Pauli group \mathcal{P}_n to be the subgroup of $\text{SU}(\mathcal{H}_n)$ generated by $\{X_j, Y_j, Z_j : j = 1, \dots, n\}$. Note that Hermiticity is inherited from the single-qubit case, i.e. $P^\dagger = P$ for every $P \in \mathcal{P}_n$.

Definition 4 (quantum one-time pad). For $r \in \{0, 1\}^{2n}$, we define the quantum one-time pad (QOTP) on n qubits with classical key r to be the map:

$$P_r := \prod_{j=1}^n X_j^{r_{2j-1}} Z_j^{r_{2j}} \in \mathcal{P}_n.$$

The effect of P_r on any quantum state $\rho \in \mathfrak{D}(\mathcal{H}_n)$ is simply

$$\frac{1}{2^{2n}} \sum_{r \in \{0, 1\}^{2n}} P_r \rho P_r = \frac{\mathbb{1}_n}{2^n}.$$

As before, the map $\rho \mapsto P_r \rho P_r$ (for uniformly random key r) is an information-theoretically secure symmetric-key encryption scheme for quantum states.

Just as in the classical case [Sha49], any reduction in key length is not possible without compromising information-theoretic security [AMTdW00, BR03]. Of course, in practice the key length of the one-time pad (quantumly or classically) is highly impractical. This is a crucial reason to consider—as we do in this work—encryption schemes which are secure only against computationally bounded adversaries.

2.3 Efficient Classical and Quantum Computations

We will refer to several different notions of efficient algorithms. The most basic of these is a deterministic polynomial-time algorithm (or PT). A PT \mathcal{A} is defined by

a polynomial-time uniform⁸ family $\mathcal{A} := \{\mathcal{A}_n\}_{n \in \mathbb{N}}$ of classical Boolean circuits over some gate set, with one circuit for each possible input size. For a bitstring x , we define $\mathcal{A}(x) := \mathcal{A}_{|x|}(x)$. We say that a function family $f : \{0, 1\}^n \rightarrow \{0, 1\}^m$ is PT-computable if there exists a PT \mathcal{A} such that $\mathcal{A}(x) = f(x)$ for all x ; it is implicit that m is a function of n which is bounded by some polynomial, e.g., the same one that bounds the running time of \mathcal{A} .

A probabilistic polynomial-time algorithm (or PPT) is again a polynomial-time uniform family of classical Boolean circuits, one for each possible input size n . The n th circuit still accepts n bits of input, but now also has an additional “coins” register of $p(n)$ input wires. Note that uniformity enforces that the function p is bounded by some polynomial. For a PPT \mathcal{A} , n -bit input x and $p(n)$ -bit coin string r , we set $\mathcal{A}(x; r) := \mathcal{A}_n(x; r)$. In contrast with the PT case, the notation $\mathcal{A}(x)$ will now refer to the random variable $\mathcal{A}(x; r)$ where $r \xleftarrow{\$} \{0, 1\}^{p(n)}$. Overloading notation slightly, $\mathcal{A}(x)$ can also mean the corresponding probability distribution; for example, the set of all possible outputs of \mathcal{A} on the input 1^n is denoted $\text{supp } \mathcal{A}(1^n)$.

We define a quantum polynomial-time algorithm (or QPT) to be a polynomial-time uniform family of quantum circuits, each composed of gates that may perform general admissible operations, chosen from some finite, universal set. A commonly-used alternative is to specify that the elements of the gate set are unitary. In terms of computational power, the models are the same [AKN98], however using admissible operations (versus unitary ones only) allows us to formalize a wider range of oracle-enabled QPT machines (see Section 2.4). In general, a QPT \mathcal{A} defines a family of admissible maps from input registers to output registers: $\mathcal{A} : \mathfrak{D}(\mathcal{H}_n) \rightarrow \mathfrak{D}(\mathcal{H}_u)$. As before, the n th circuit in the family will be denoted by \mathcal{A}_n . When ρ is an n -qubit state, $\mathcal{A}(\rho)$ denotes the corresponding $u(n)$ -qubit output state (by uniformity, u is bounded by some polynomial). Overloading the notation even further, for n -bit strings x we set $\mathcal{A}(x) := \mathcal{A}(|x\rangle\langle x|)$. The expression $\mathcal{A}(x) = y$ for classical y is taken to evaluate to true if the output register of the circuit contains the state $|y\rangle\langle y|$ exactly. Unless explicitly stated, any statements about the probability of an event involving a QPT are taken over the measurements of the QPT, in addition to any indicated random variables. For instance, the expression $\Pr_{x \in_R \{0, 1\}^n} [\mathcal{A}(x) = y]$ means the probability that, given a uniformly random input string x , the output register of the n th circuit of the QPT \mathcal{A} executed on $|x\rangle\langle x|$, after all gates and measurements have been applied, is in the state $|y\rangle\langle y|$.

At times, we will define QPTs with many input and output quantum registers. In these cases, some straightforward bookkeeping (e.g., via an additional classical register) may be required; for the sake of clarity, we will simply assume that this has been handled.

Throughout this work, we are concerned only with polynomial-time *uniform* computation. That is to say, the circuit families that describe any PT, PPT, or

⁸ Recall that polynomial-time uniformity means that there exists a polynomial-time Turing machine which, on input n in unary, prints a description of the n th circuit in the family.

QPT will always be both of polynomial length *and* generatable by some fixed (classical) polynomial-time Turing machine. In particular, we consider uniform adversaries only—although all of our results carry over appropriately to the non-uniform setting as well.

2.4 Oracles

We denote by \mathcal{A}^f an algorithm which has oracle access to some function family f . Such an algorithm (whether PT, PPT, or QPT) is defined as above, except each circuit in the algorithm can make use of additional “oracle gates” (one for each possible input size) which evaluate f . In the case of PTs and PPTs, oracles can implement any function from bitstrings to bitstrings. In the case of QPTs, we consider two different oracle types.

First, we allow purely classical oracles. Just as in the case of PTs and PPTs, a classical oracle implements a function f from bitstrings to bitstrings. In the case of a QPT with a classical oracle, *queries can be made on classical inputs only* (this is sometimes referred to as “standard-security” [Zha12]). We emphasize that we do not require that the oracle is made reversible, nor do we allow the QPT to input superpositions. Note that any such oracle can be implemented by an admissible map, such that classical inputs x are deterministically mapped to $f(x)$ (to see this, start with a Boolean circuit for f , make it reversible, and then recall that adding ancillas and discarding output bits are admissible operations). While it might seem that disallowing superposition inputs is an artificial and unrealistic restriction, in our case it actually strengthens results. For instance, we will show that secure quantum encryption can be achieved using pseudorandom functions which are secure only against quantum adversaries possessing just classical oracle access. One can of course also ask for *more powerful* functions (which are secure against superposition access, or “quantum-secure” [Zha12]) but this turns out to be unnecessary in our case. Second, we also allow oracles that are admissible maps. More precisely, for an admissible map family \mathcal{C} , we write $\mathcal{A}^{\mathcal{C}}$ to denote a QPT whose circuits can make use of special “oracle gates” which implement admissible maps from the family \mathcal{C} . Each such gate accepts a quantum register as input, to which it applies the appropriate admissible map from the family, and returns an output register. It is not necessary for the input and output registers to have the same number of qubits.

In any case, each use of an oracle gate counts towards the circuit length, and hence also towards the total computation time of the algorithm. In particular, no PT, PPT or QPT algorithm may make more than a polynomial number of oracle calls.

3 Quantum Encryption and Indistinguishability

In this section, we give general definitions of encryption schemes for quantum data (Section 3.1) and a corresponding notion of indistinguishability, including IND-CPA and IND-CCA1 (Section 3.2.)

3.1 Quantum Encryption Schemes

We start by defining *secret-key encryption for quantum data*. In the following we assume that the secret key is a classical bitstring, while the plaintext and the ciphertext can be arbitrary quantum states. We refer to \mathcal{K} , \mathcal{H}_M and \mathcal{H}_C as the key space, the message (or plaintext) space, and the ciphertext space, respectively. We remark that these are actually infinite families of spaces, each with a number of (qu)bits which scales polynomially with n . We assume that $\mathcal{K} := \{0, 1\}^n$, so that the key-length is n bits, and the plaintext and the ciphertext lengths are $m \leq \text{poly}(n)$ and $c \leq \text{poly}(n)$ qubits, respectively. The key-generation algorithm accepts a description of the security parameter n in unary and outputs a classical key of length n . Later, we will define an additional Hilbert space \mathcal{H}_E in order to model auxiliary information used by some adversary. Encryption accepts a classical key and a plaintext, and outputs a ciphertext; decryption accepts a classical key and a ciphertext, and outputs a plaintext. The correctness guarantee is that plaintexts are preserved (up to negligible error) under encryption followed by decryption under the same key.

Definition 5. A quantum symmetric-key encryption scheme (or qSKE) is a triple of QPTs:

1. (key generation) $\text{KeyGen} : 1^n \mapsto k \in \mathcal{K}$
2. (encryption) $\text{Enc} : \mathcal{K} \times \mathfrak{D}(\mathcal{H}_M) \rightarrow \mathfrak{D}(\mathcal{H}_C)$
3. (decryption) $\text{Dec} : \mathcal{K} \times \mathfrak{D}(\mathcal{H}_C) \rightarrow \mathfrak{D}(\mathcal{H}_M)$

such that $\|\text{Dec}_k \circ \text{Enc}_k - \mathbb{1}_M\|_\diamond \leq \text{negl}(n)$ for all $k \in \text{supp KeyGen}(1^n)$.

In the above, we used a convenient shorthand notation for encryption and decryption maps with a fixed key k (which is classical), formally defined by $\text{Enc}_k : \rho \mapsto \text{Enc}(k, \rho)$ and $\text{Dec}_k : \sigma \mapsto \text{Dec}(k, \sigma)$.

Next, we define a notion of *public-key encryption for quantum data*. In addition to the usual spaces from the symmetric-key setting above, we now also have a public key of length $p(n) \leq \text{poly}(n)$ bits. We define the related public-key space as $\mathcal{K}_{pub} \subset \{0, 1\}^p$ and reuse \mathcal{K} for the corresponding private-key space.

Definition 6. A quantum public-key encryption scheme (or qPKE) is a triple of QPTs:

1. (key-pair generation) $\text{KeyGen} : 1^n \mapsto (pk, sk) \in \mathcal{K}_{pub} \times \mathcal{K}$
2. (encryption with public key) $\text{Enc} : \mathcal{K}_{pub} \times \mathfrak{D}(\mathcal{H}_M) \rightarrow \mathfrak{D}(\mathcal{H}_C)$
3. (decryption with private key) $\text{Dec} : \mathcal{K} \times \mathfrak{D}(\mathcal{H}_C) \rightarrow \mathfrak{D}(\mathcal{H}_M)$

such that $\|\text{Dec}_{sk} \circ \text{Enc}_{pk} - \mathbb{1}_m\|_\diamond \leq \text{negl}(n)$ for all $(pk, sk) \in \text{supp KeyGen}(1^n)$.

In this case, we again placed the relevant keys in the subscript, i.e.,

$$\text{Enc}_{pk} : \rho \mapsto \text{Enc}(pk, \rho) \quad \text{and} \quad \text{Dec}_{sk} : \sigma \mapsto \text{Dec}(sk, \sigma).$$

We remark that some variations of the above two definitions are possible. For instance, one could demand that encryption followed by decryption is exactly equal to the identity operator. The schemes we present in [Section 5](#) will in fact satisfy this stronger condition.

3.2 Indistinguishability of Encryptions

Following the classical definition, the security notion of *quantum indistinguishability under chosen plaintext attacks* has been considered previously for the case of quantum encryption schemes in [BJ15] and for classical encryption schemes in [GHS15]. Here, we present the definition from [BJ15], which we slightly extend to the CCA1 (chosen ciphertext attack) setting. The security definitions are formulated with the public-key (or asymmetric-key) setting in mind, and we clarify when meaningful differences in the symmetric-key setting arise.

Our definition models a situation in which an honest user encrypts messages of the adversary's choice; the adversary then attempts to match the ciphertexts to the plaintexts. In our formulation, an IND adversary consists of two QPTs: the *message generator* and the *distinguisher*. The message generator takes as input the security parameter and a public key, and outputs a challenge state consisting of a plaintext and some auxiliary information. The auxiliary information models, for instance, the fact that the output state might be entangled with some internal state of the adversary itself. Then the distinguisher receives this auxiliary information, and a state which might be either the encryption of the original challenge state or the encryption of the zero state. The distinguisher's goal is to decide which of the two is the case.

Security in this model requires that the adversary does not succeed with probability significantly better than guessing. We also define two standard variants: indistinguishability under chosen plaintext attack (IND-CPA) and indistinguishability under chosen-ciphertext-attack (IND-CCA1). We leave the definition of CCA2 (adaptive chosen ciphertext attack) security as an interesting open problem. As before, all circuits are indexed by the security parameter.

Definition 7 (IND). *A qPKE scheme (KeyGen, Enc, Dec) has indistinguishable encryptions (or is IND secure) if for every QPT adversary $\mathcal{A} = (\mathcal{M}, \mathcal{D})$ we have:*

$$\left| \Pr \left[\mathcal{D} \{ (\text{Enc}_{pk} \otimes \mathbb{1}_E) \rho_{ME} \} = 1 \right] - \Pr \left[\mathcal{D} \{ (\text{Enc}_{pk} \otimes \mathbb{1}_E) (|0\rangle \langle 0|_M \otimes \rho_E) \} = 1 \right] \right| \leq \text{negl}(n)$$

where $\rho_{ME} \leftarrow \mathcal{M}(pk)$, $\rho_E = \text{Tr}_M(\rho_{ME})$, and the probabilities are taken over $(pk, sk) \leftarrow \text{KeyGen}(1^n)$ and the internal randomness of Enc, \mathcal{M} , and \mathcal{D} .

- **IND-CPA:** In addition to the above, \mathcal{M} and \mathcal{D} are given oracle access to Enc_{pk} .
- **IND-CCA1:** In addition to IND-CPA, \mathcal{M} is given oracle access to Dec_{sk} .

Here we use $|0\rangle \langle 0|_M$ to denote $|0^m\rangle \langle 0^m|$, where m is the number of qubits in the M register.

The definition is illustrated in Figure 1. The symmetric-key scenario is the same, except $pk = sk$, and \mathcal{M} receives only a blank input. We remark that in the public-key setting, IND implies IND-CPA: an adversary with knowledge of pk can easily simulate the Enc_{pk} oracle. Note that, under CPA, the IND definition is known to be equivalent to IND in the *multiple-message* scenario [BJ15].




Fig. 1. IND posits that a QPT $(\mathcal{M}, \mathcal{D})$ cannot distinguish between these two scenarios.

4 Quantum Semantic Security

This section is devoted to defining quantum semantic security (Section 4.2), and showing its equivalence with quantum indistinguishability (Section 4.3).

Following the classical definition, the security notion of *quantum semantic security under chosen plaintext attacks* has been given previously in [GHS15] for the case of a special class of quantum states arising when considering quantum access to classical encryption schemes. Here, we give a more general definition for arbitrary quantum plaintexts. As we outlined the classical situation with semantic security in Section 1.1.1, we start with a discussion of some difficulties in transitioning to the quantum setting. A similar discussion can be found in [GHS15] and we explain below where and why we make different choices.

4.1 Difficulties in the Quantum Setting

When attempting to transfer the definition of semantic security to the quantum world, the main question one encounters is to determine the quantum equivalents of $h(x)$ and $f(x)$ (as it is relatively clear that the plaintext x would have as quantum equivalent a quantum state ρ_M , in a *message register*, M).

For the case of the side-information, $h(x)$, one might attempt to postulate that this side information is available via the output of a quantum map Φ_h , evaluated on ρ_M . There are, however, two obvious problems with this approach: firstly, it is unclear how to *simultaneously* generate both ρ_M and $\Phi_h(\rho_M)$ (the main obstacle stemming from the quantum *no-cloning* theorem [WZ82], according to which it is not possible to perfectly copy an unknown quantum state)⁹. Secondly, it is well-established that the most general type of quantum side-information includes entanglement (contrary to the scenario studied in [GHS15]). We therefore conclude that side information should be modelled simply as an extra register (called E) such that ρ_{ME} are in an arbitrary quantum state (as generated by some process—for a formal description, see Definition 8).

⁹ [GHS15] solves the issue by requiring a quantum circuit that takes classical randomness as input and outputs plaintext states. Hence, multiple plaintext states can be generated by using the same randomness.

For the case of the target function f , one might also postulate a quantum map Φ_f , the goal then (for both the adversary and simulator), being to output $\Phi_f(\rho_M)$. However, given that quantum states and maps form a continuum, one must exercise care in quantifying when a simulator has successfully simulated the adversary. We propose three possible tests for quantifying “success” in the semantic security game, each leading to its own definition. Since we show that all three definitions are equivalent, we conclude that it is a matter of taste (or context) which definition to label as *the* definition of quantum semantic security. We focus in this section on the first one, which we called SEM, because we find that it the most natural. We give formal definitions and proofs of equivalence for all three definitions in Appendix A. Here is an overview of the three different notions:

- **SEM.** In Definition 8, a state ρ_{MEF} is generated; intuitively, the contents of register F can be seen as a “target” output that the adversary tries to achieve (however, this is not quite the case as we point out shortly). We then postulate a quantum polynomial time *distinguisher* who is given the F register and charged with distinguishing the output of the adversary from the output of the simulator, with security being associated with the inability of the distinguisher in telling the two situations apart. We thus see that the role of register F is actually to assist the distinguisher: semantic security corresponds to the situation where the distinguisher essentially cannot tell the real from ideal apart, *even with access to the F system*.
- **SEM2.** In Definition 22, we specify instead that the state ρ_{MEF} be a *classical-quantum state*. That is, ρ_{ME} is quantum, but the register F contains a classical state. Thus, correlations shared between the two systems are classical only. The requirement for security is that the simulator should provide a classical output that equals the contents of F , essentially just as well as the adversary can.
- **SEM3.** In Definition 25, we introduce a classical function f , thus closely mimicking the classical definition. Namely, we specify as in SEM2 that F contains a classical state y , which we furthermore assume to be precisely the results of any measurements used to generate ρ_{ME} (thus, y is, in a sense, a full “classical description” of ρ_{ME}). The requirement for security is that the simulator is able to output $f(y)$ (for any f) with essentially the same probability as the adversary.

4.2 Definition of Semantic Security

As before, we work primarily in the public-key setting; adaptation to the symmetric-key setting is again straightforward. In our concrete formulation of SEM (Definition 8), we define the following QPT machines: the *message generator* \mathcal{M} (which generates ρ_{MEF}), the *adversary* \mathcal{A} , the *simulator* \mathcal{S} and the *distinguisher* \mathcal{D} .

Definition 8. [SEM] A qPKE scheme (KeyGen, Enc, Dec) is semantically secure if for any QPT adversary \mathcal{A} , there exists a QPT simulator \mathcal{S} such that for all

QPTs \mathcal{M} and \mathcal{D} ,

$$|\Pr [\mathcal{D}\{(\mathcal{A} \otimes \mathbb{1}_F)(\text{Enc}_{pk} \otimes \mathbb{1}_{EF})\rho_{MEF}\} = 1] - \Pr [\mathcal{D}\{(\mathcal{S} \otimes \mathbb{1}_F)\rho_{EF}\} = 1]| \leq \text{negl}(n),$$

where $\rho_{MEF} \leftarrow \mathcal{M}(pk)$, $\rho_{EF} = \text{Tr}_M(\rho_{MEF})$, and the probability is taken over $(pk, sk) \leftarrow \text{KeyGen}(1^n)$ and the internal randomness of Enc , \mathcal{A} , \mathcal{S} and \mathcal{D} .

- **SEM-CPA:** In addition to the above, all QPTs are given oracle access to Enc_{pk} .
- **SEM-CCA1:** In addition to IND-CPA, \mathcal{M} is given oracle access to Dec_{sk} .

The interactions among the QPTs are illustrated in [Figure 2](#). A few remarks are in order. First, all the registers above are uniformly of size polynomial in n . Second, the input and output registers of the relevant QPTs are understood from context, e.g., the expression $(\mathcal{S} \otimes \mathbb{1}_F)\rho_{EF}$ makes clear that the input register of \mathcal{S} is E . Third, we note that SEM implies SEM-CPA in the public-key setting, since access to the public key implies simulatability of Enc_{pk} . Finally, just as in the case of IND, adapting to the symmetric-key setting is simply a matter of setting $pk = sk$ and positing that \mathcal{M} receives only a blank input.




Fig. 2. SEM: for all adversaries \mathcal{A} there exists a simulator \mathcal{S} such that these two scenarios are indistinguishable.

The classical (uniform) definition of semantic security is recovered as a special case, as follows. All of the QPTs are PPTs, and the message generator \mathcal{M} outputs classical plaintext m , side information $h(m)$ and target function $f(m)$. The distinguisher \mathcal{D} simply checks whether the adversary's (or simulator's) output is equal to the contents of the F register.

4.3 Semantic Security is Equivalent to Indistinguishability

While semantic security gives a strong and intuitively meaningful definition of security, indistinguishability is typically easier to prove and work with. In this section we show that—just as in the classical setting—the two notions are equivalent. This proves [Theorem 1](#). The equivalence holds for all of the variants of

Definition 7 and **Definition 8**: under either public or private-key, we have equivalence of IND with SEM, IND-CPA with SEM-CPA, and IND-CCA1 with SEM-CCA1. Here, we focus on the SEM definition; see **Appendix A** for the equivalence with the SEM2 and SEM3 definitions.

Theorem 9 (IND \implies SEM). *If a quantum encryption scheme $(\text{KeyGen}, \text{Enc}, \text{Dec})$ has indistinguishable encryptions (IND), then it is semantically secure (SEM).*

Proof. Suppose that an encryption scheme $(\text{KeyGen}, \text{Enc}, \text{Dec})$ has indistinguishable encryptions. Let \mathcal{A} be QPT SEM attacker against semantic security as in **Definition 8**. We define the QPT SEM simulator \mathcal{S} as follows: \mathcal{S} does not receive $\text{Enc}_{pk}(\rho_M)$, but instead runs \mathcal{A} on input $(\text{Enc}_{pk} \otimes \mathbb{1}_E)(|0\rangle\langle 0| \otimes \rho_E)$ and outputs whatever \mathcal{A} outputs. Let \mathcal{M} be a QPT SEM message generator that outputs ρ_{MEF} .

Assume for a contradiction the existence of a QPT SEM distinguisher \mathcal{D} which successfully distinguishes the output of \mathcal{A} from the output of \mathcal{S} (with the help of register F), then the combination of \mathcal{A} and \mathcal{D} successfully distinguishes $(\text{Enc}_{pk} \otimes I_{EF})\rho_{MEF}$ from $(\text{Enc}_{pk} \otimes I_{EF})(|0\rangle\langle 0| \otimes \rho_{EF})$, hence contradicting the indistinguishability. \square

In the private-key setting without CPA oracle access, \mathcal{S} runs $\text{KeyGen}(1^n)$ to generate his own secret key k' , and then encrypts $|0^n\rangle\langle 0^n|$ using k' instead of k . The ciphertexts $\text{Enc}_k |0\rangle\langle 0|$ and $\text{Enc}_{k'} |0\rangle\langle 0|$ will be distributed identically since k and k' are. Hence, the success probability of the SEM simulator \mathcal{S} does not change.

In case of CPA and CCA1 oracles, both for the public- and private-key setting, the simulator \mathcal{S} forwards \mathcal{A} 's oracle queries to his own oracle(s), and \mathcal{S} obtains \mathcal{A} 's input state by a call to his encryption oracle on state $|0\rangle\langle 0|$, joined with his auxiliary information ρ_E .

Theorem 10 (SEM \implies IND). *If a quantum encryption scheme $(\text{KeyGen}, \text{Enc}, \text{Dec})$ is semantically secure (SEM), then it has indistinguishable encryptions (IND).*

Proof. Let $(\mathcal{M}, \mathcal{D})$ be an IND adversary such that \mathcal{D} distinguishes $(\text{Enc}_{pk} \otimes \mathbb{1}_E)\rho_{ME}$ from $(\text{Enc}_{pk} \otimes \mathbb{1}_E)(|0\rangle\langle 0| \otimes \rho_E)$ with advantage $\varepsilon(n)$ if $\rho_{ME} \leftarrow \mathcal{M}$. Let us consider the SEM message generator \mathcal{M}' which runs $\rho_{ME} \leftarrow \mathcal{M}$ and outputs (with probability $\frac{1}{2}$ each) either the state $\rho_{ME} \otimes |0\rangle\langle 0|_F$ or the state $|0\rangle\langle 0|_M \otimes \rho_E \otimes |1\rangle\langle 1|_F$. Next we consider the SEM attacker \mathcal{A} which runs \mathcal{D} and outputs the classical bit that \mathcal{D} outputs. We also consider the SEM attacker $\mathcal{A} \oplus 1$, which outputs the opposite bit. As SEM distinguisher, let us consider the procedure which compares \mathcal{A} 's output bit to a measurement (in the computational basis) of the qubit in register F . Any SEM simulator \mathcal{S} that does not have access to the encrypted M -register has to guess the state of the random bit in F and will be correct with probability $1/2$. Then $\varepsilon(n)$ is twice the maximum of the advantages that \mathcal{A} and $\mathcal{A} \oplus 1$ have in successfully predicting F over $1/2$, the probability of success of any simulator. By SEM, both of these advantages are negligible, and hence so is $\varepsilon(n)$. \square

5 Quantum Encryption Schemes

We now turn to the question of existence for encryption schemes for quantum data. We present two schemes based on the existence of classical functions which are difficult to invert for quantum computers. The first scheme (Section 5.1) is symmetric-key and IND-CCA1-secure; the second scheme (Section 5.2) is public-key and IND-CPA-secure. By the results of Section 4, these schemes are also semantically secure.

5.1 Quantum Symmetric-Key Encryption from One-Way Functions

In this section, we prove [Theorem 2](#): *If quantum-secure one-way functions exist, then so do IND-CCA1-secure private-key quantum encryption schemes.*

The proof proceeds in two steps. First, we define quantum-secure one-way functions (qOWFs) and quantum-secure pseudo-random functions (qPRFs); we can argue as in the classical world that qPRFs exist if qOWFs do ([Theorem 13](#).) Second, we show that any qPRF can be used to construct an explicit IND-CCA1-secure symmetric-key scheme for quantum data.

We begin with the formal definitions of qOWFs and qPRFs, and a statement of the result connecting the two.

Definition 11. *A PT-computable function $f : \{0, 1\}^* \rightarrow \{0, 1\}^*$ is a quantum-secure one-way function (qOWF) if for every QPT \mathcal{A} ,*

$$\Pr_{x \xleftarrow{\$} \{0,1\}^n} [\mathcal{A}(f(x), 1^n) \in f^{-1}(f(x))] \leq \text{negl}(n).$$

Definition 12. *A PT-computable function family $f : \{0, 1\}^n \times \{0, 1\}^m \rightarrow \{0, 1\}^\ell$ is a quantum-secure pseudorandom function (qPRF) if for every QPT \mathcal{D} equipped with a classical oracle,*

$$\left| \Pr_{k \xleftarrow{\$} \{0,1\}^n} [\mathcal{D}^{f_k}(1^n) = 1] - \Pr_{g \xleftarrow{\$} \{\{0,1\}^m \rightarrow \{0,1\}^\ell\}} [\mathcal{D}^g(1^n) = 1] \right| \leq \text{negl}(n).$$

We remark that, to some readers, the restriction to classical oracles might seem artificial. While one can certainly consider functions with the *stronger* guarantee of resistance to quantum adversaries with quantum oracle access, stronger functions are not necessary to establish our results. We thus opt for the weaker primitive. In either case, the following holds.

Theorem 13. *If qOWFs exist, then qPRFs exist.*

Since our definitions are in terms of *classical* oracles, the classical proof that shows that qOWFs imply qPRFs carries through [[HILL99](#), [GGM86](#)]. We remark that Zhandry [[Zha12](#)] extended this result to the case of functions secure against quantum superposition queries, what he calls “quantum-secure PRFs.” It should

be noted that the proof of the Theorem 13 actually implies the existence of a qPRF for any (polynomial) choice of the parameters m and ℓ in Definition 12.

We are now ready to proceed with the second part of the proof of Theorem 2, namely the construction of an encryption scheme from a given qPRF. Essentially, this scheme encrypts a quantum state ρ by first selecting a random string r , then inputting r into a qPRF; the output $f_{k(r)}$ is then used as an encryption key for the quantum one-time pad, $P_{f_{k(r)}}$.

Scheme 1 Let $f : \{0, 1\}^n \times \{0, 1\}^{2n} \rightarrow \{0, 1\}^{2n}$ be a qPRF. Let qPRF-SKE be the following triple of QPT algorithms:

1. (key generation) $\text{KeyGen}(1^n)$: output $k \xleftarrow{\$} \{0, 1\}^n$;
2. (encryption) $\text{Enc}_k(\rho)$: choose $r \xleftarrow{\$} \{0, 1\}^{2n}$ and output $|r\rangle\langle r| \otimes P_{f_k(r)}\rho P_{f_k(r)}$.
3. (decryption) $\text{Dec}_k(\sigma)$: measure the first $2n$ qubits in the computational basis to obtain $r' \in \{0, 1\}^{2n}$; apply $P_{f_k(r')}$ to remaining $2n$ qubits and output the result.

For simplicity, we chose $\mathfrak{D}(\mathcal{H}_n)$ for the key space and the plaintext space, and $\mathfrak{D}(\mathcal{H}_{2n})$ for the ciphertext space; we can easily adapt the above to other polynomially-related cases by selecting a qPRF with different parameters. Correctness of Scheme 1 is easily verified:

$$\text{Dec}_k(\text{Enc}_k(\rho)) = \text{Dec}_k(|r\rangle\langle r| \otimes P_{f_k(r)}\rho P_{f_k(r)}) = P_{f_k(r)}P_{f_k(r)}\rho P_{f_k(r)}P_{f_k(r)} = \rho,$$

where the second equality follows from the definition of the decryption function and the last step is due to the fact that the Pauli operators are self-inverse. Next, we show that the scheme is secure against non-adaptive chosen ciphertext attacks. The classical version of this result is standard, and we use essentially the same proof; see, e.g., Proposition 5.4.18 in Goldreich's textbook [Gol04b].

Lemma 14. *If f is a qPRF, then Scheme 1 is an IND-CCA1-secure symmetric-key quantum encryption scheme as defined in Definition 7.*

Proof. First, we analyse the security of the scheme in an idealized scenario where f is a truly random function. We claim that in this case, \mathcal{A} correctly guesses the challenge with probability at most $1/2 + \text{negl}(n)$ (see Definition 27). In fact, this bound holds for a stronger adversary \mathcal{A}' , who has access to a classical oracle for f prior to the challenge, and access to polynomially-many pairs $(r_i, f(r_i))$ for random $r_i, 1 \leq i \leq q$, after the challenge. This adversary is stronger than \mathcal{A} since it can simulate \mathcal{A} by implementing Enc_f and Dec_f oracles using its f oracles. Since the input r into f in the challenge ciphertext is uniformly random, the probability that any of the polynomially-many oracle calls of \mathcal{A}' uses the same r is negligible. In the case that no oracle calls use r , the mixtures of the inputs to \mathcal{A}' (including the pairs $(r_i, f(r_i))$) are the same for the original challenge and the zero challenge. This fact can be verified by first averaging over the values of $f(r)$: since f is uniformly random, $f(r)$ is also uniformly random as well

as independent of the other values of f . In both cases, applying the quantum one-time pad results in the state:

$$|r\rangle \langle r| \otimes \frac{1}{2^n} \mathbb{1} \otimes \rho_E \otimes |r_1\rangle \langle r_1| \otimes |f(r_1)\rangle \langle f(r_1)| \otimes \cdots \otimes |r_q\rangle \langle r_q| \otimes |f(r_q)\rangle \langle f(r_q)|,$$

and indistinguishability follows.

Next, we consider the case that f is a pseudorandom function. We show that a successful IND-CCA1 adversary \mathcal{A} (i.e., one that distinguishes challenges with better than negligible probability) can be used to construct a successful f -adversary \mathcal{A}_0 (i.e., one that distinguishes f from random with non-negligible probability.) The adversary \mathcal{A}_0 is a QPT with classical oracle access to a function $\varphi : \{0, 1\}^{2n} \rightarrow \{0, 1\}^{2n}$, and aims to output 0 if φ is perfectly random and 1 if $\varphi = f_k$ for some k . Define the simulated oracles

$$\text{Enc}_\varphi : \rho \mapsto \left(r, P_{\varphi(r)} \rho P_{\varphi(r)} \right) \text{ for } r \xleftarrow{\$} \{0, 1\}^{2n} \quad \text{and} \quad \text{Dec}_\varphi : |r'\rangle \langle r'| \otimes \rho \mapsto P_{\varphi(r')} \rho P_{\varphi(r')},$$

where, as before, we assume that Dec_φ measures the first register before decrypting the second. Note that if $\varphi = f_k$ then these are exactly the encryption and decryption oracles (with key k) of the qPRF-SKE scheme.

The QPT \mathcal{A}_0^φ proceeds as follows. First, it simulates \mathcal{A} , and replies to its queries to the encryption oracle with Enc_φ and its queries to the decryption oracle with Dec_φ . When it transmits the challenge, \mathcal{A}_0^φ replies with either the encryption of the challenge, or the encryption of $|0^n\rangle \langle 0^n|$, each with probability $1/2$. If \mathcal{A} responds correctly, \mathcal{A}_0^φ outputs 1; otherwise it outputs 0. If $\varphi = f_k$ then we have exactly simulated the IND-CCA1 game with adversary \mathcal{A} ; in that case, since \mathcal{A} is IND-CCA1-breaking, \mathcal{A}_0^φ outputs 1 with probability at least $1/2 + 1/p(n)$ for some polynomial p , for infinitely many n .

We conclude that

$$\left| \Pr_{k \xleftarrow{\$} \{0, 1\}^n} [\mathcal{A}_0^{f_k}(1^n) = 1] - \Pr_{\varphi \xleftarrow{\$} \{\{0, 1\}^{2n} \rightarrow \{0, 1\}^{2n}\}} [\mathcal{A}_0^\varphi(1^n) = 1] \right| \geq 1/p(n) - \text{negl}(n),$$

for infinitely many n , i.e., f is not a qPRF. \square

Putting together [Theorem 13](#) and [Lemma 14](#), we arrive at a proof of [Theorem 2](#).

5.2 Quantum Public-Key Encryption from Trapdoor Permutations

For the construction of public-key schemes, we will need qOWFs with an additional property: the existence of *trapdoors* which enable efficient inversion. Following the classical approach of Diffie and Hellman [[DH76](#)], we set down the notion of a quantum-secure trapdoor one-way permutation (or qTOWP), and then show how to use any qTOWP to construct IND-CPA secure public-key encryption schemes for quantum data. This will establish [Theorem 3](#): *If*

quantum-secure trapdoor one-way permutations exist, then so do semantically secure public-key quantum encryption schemes..

We begin with a definition of qTOWPs. We require a slight (but standard) variation of [Definition 11](#), namely the notion of a quantum-secure one-way permutation (or qOWP). A qOWP is a qOWF whose input domains are sets D_i ; moreover, the function restricted to any such domain must be a permutation (from the domain to the corresponding range.) When we augment such a qOWP with trapdoors, we arrive at the following definition.

Definition 15. A quantum-secure trapdoor one-way permutation (qTOWP) is a qOWF

$$\{f_i : D_i \rightarrow \{0, 1\}^*\}_{i \in I}$$

(where each f_i is a bijection), together with a triple of PPTs $(\mathcal{G}, \mathcal{S}, \mathcal{I})$ which

1. (generate (index, trapdoor) pair) $\text{supp } \mathcal{G}(1^n) \subseteq (I \cap \{0, 1\}^n) \times \{0, 1\}^n$;
2. (sample from domain) for all $i \in I$, $\text{supp } \mathcal{S}(i) = D_i$;
3. (invert using trapdoor) for all $(i, t) \in \text{supp } \mathcal{G}(1^n)$ and all $x \in D_i$, $\mathcal{I}(f_i(x), t) = x$.

Before we can describe the public-key scheme and prove its security, we need two additional (well-known) primitives which can be constructed from any qOWP, with or without trapdoors. The first is a quantum-secure “hard-core” predicate, which is a “yes” or “no” question about inputs x which is difficult to answer if one only knows $f(x)$.

Definition 16. A PT-computable $b : \{0, 1\}^* \rightarrow \{0, 1\}$ is a hard-core of a qOWP f if for every QPT \mathcal{A} ,

$$\Pr_{x \xleftarrow{\$} \{0, 1\}^n} [\mathcal{A}(f(x), 1^n) = b(x)] \leq \frac{1}{2} + \text{negl}(n).$$

Theorem 17. ([\[AC02\]](#), quantum analogue of [\[GL89\]](#)) If qOWPs exist, then qOWPs with hard-cores exist.

The other primitive we need is a quantum-secure pseudorandom generator, which is defined below. The classical proof that hard-cores imply pseudorandom generators carries over with little modification (see [Lemma 19](#)).

Definition 18. A PT-computable deterministic function $G : \{0, 1\}^n \rightarrow \{0, 1\}^m$ is a quantum-secure pseudorandom generator (qPRG) if for every QPT \mathcal{D} ,

$$\left| \Pr_{s \xleftarrow{\$} \{0, 1\}^n} [\mathcal{D}(G(s)) = 1] - \Pr_{y \xleftarrow{\$} \{0, 1\}^m} [\mathcal{D}(y) = 1] \right| \leq \text{negl}(n).$$

Lemma 19. Suppose f is a qOWP, b its hard-core predicate, and let t be polynomial in n . Then $G : s \mapsto b(f^{t-1}(s))b(f^{t-2}(s)) \dots b(s)$ is a qPRG.

Proof (Sketch). The proof proceeds almost identically as in the classical case (see, e.g., [Gol04a].) Let \mathcal{D} be a quantum adversary that distinguishes $G(U_n)$ from uniform. Note that, as stated in Definition 18, \mathcal{D} gets only classical bitstring outputs from the pseudorandom generator. In the classical proof, one constructs an adversary \mathcal{A} which uses \mathcal{D} as a black-box subroutine, and breaks the hard-core of f . We use the exact same \mathcal{A} now; in particular, we only need to invoke \mathcal{D} on classical inputs and read out its (post-measurement) classical outputs (0 or 1). Of course, by virtue of needing to invoke \mathcal{D} , \mathcal{A} itself will now be a QPT.

In slightly greater detail, we use a standard hybrid argument to give a “predictor” algorithm \mathcal{A} that, for some index $i \leq t$, can predict the $i + 1^{\text{st}}$ bit of $G(U_n)$, given as input the first i bits of the output of G . \mathcal{A} succeeds with non-negligible advantage over random, i.e., the probability over s that $\mathcal{A}(b(f^{t-1}(s)) \dots b(f^{t-i}(s)))$ outputs $b(f^{t-(i+1)}(s))$ is at least $1/2 + 1/p(n)$ where $p(n)$ is some polynomial. Crucially, since f implements a permutation over $\{0, 1\}^n$, we have that $b(f^{i-1}(U_n)) \dots b(U_n)$ is distributed identically to $b(f^{t-1}(U_n)) \dots b(f^{t-i}(U_n))$. Therefore, given uniform x , and $y = f(x)$, we can use the output of the predictor, $A(b(f^{i-1}(y)) \dots b(y)) = A(b(f^i(x)) \dots b(f(x)))$ to predict $b(x)$ with non-negligible advantage, in violation of the security guarantee of the hard-core predicate. \square

We now have all of the ingredients needed to describe a public-key scheme for encrypting quantum data.

Scheme 2 Let f be a qTOWP, and let b and $G : \{0, 1\}^n \rightarrow \{0, 1\}^{2n}$ be a corresponding hard-core and qPRG, respectively. Let qTOWP-PKE be the following triple of algorithms:

1. ((public, private) key-pair generation) $\text{KeyGen}(1^n)$: output $\mathcal{G}(1^n) = (i, t) \in \{0, 1\}^n \times \{0, 1\}^n$;
2. (encryption with public key) $\text{Enc}_i(\rho)$:
 - apply $\mathcal{S}(i)$ to select $d \in D_i$, and compute $r := G(d)$;
 - output $|f_i^{2n}(d)\rangle \langle f_i^{2n}(d)| \otimes P_r \rho P_r$
3. (decryption with private key) $\text{Dec}_t(|s\rangle \langle s| \otimes \sigma)$:
 - for $j = 1, \dots, 2n$, apply $b \circ (\mathcal{I})^j$ to (s, t) ; concatenate the resulting bits to get $u \in \{0, 1\}^{2n}$;
 - output $P_u \sigma P_u$.

Correctness of the scheme is straightforward; fix a key-pair (i, t) , a randomly sampled $d \in D_i$, and the corresponding r . Then

$$\text{Dec}_t(\text{Enc}_i(\rho)) = \text{Dec}_t(|f_i^{2n}(d)\rangle \langle f_i^{2n}(d)| \otimes P_r \rho P_r) = P_u P_r \rho P_r P_u = \rho,$$

where the last step follows from the fact that $u = r$ for valid ciphertexts. It remains to show that this scheme is secure against chosen-plaintext attacks. We begin by proving indistinguishability of ciphertexts for the quantum one-time pad which uses randomness supplied by a qPRG. We first set the following notation. Recall from Section 2.2 that a string r of $2n$ bits determines a Pauli group element $P_r \in U(2^n)$. Given an n -qubit register A , an arbitrary register B , and $\rho \in \mathfrak{D}(\mathcal{H}_A \otimes H_B)$, define $\mathbb{P}_{r:A}(\rho) := (P_r \otimes \mathbb{1}_B)\rho(P_r \otimes \mathbb{1}_B)$.

Lemma 20. *Suppose $G : \{0, 1\}^n \rightarrow \{0, 1\}^m$ is a qPRG. Then for any efficiently preparable states $\rho_{AB} \in \mathfrak{D}(\mathcal{H}_A \otimes \mathcal{H}_B)$ and $\sigma_A \in \mathfrak{D}(\mathcal{H}_A)$, and any QPT \mathcal{D} ,*

$$\left| \Pr_{s \xleftarrow{\$} \{0,1\}^n} \left[\mathcal{D}(\mathbb{P}_{G(s);A}(\rho_{AB})) = 1 \right] - \Pr_{s \xleftarrow{\$} \{0,1\}^n} \left[\mathcal{D}(\mathbb{P}_{G(s);A}(\sigma_A \otimes \rho_B)) = 1 \right] \right| \leq \text{negl}(n). \quad (1)$$

Proof. The two key observations are (i.) distinguishability as in Equation (1) is impossible if we replace $G(s)$ with uniform randomness, and (ii.) with only classical input/output access to G , we can simulate $\mathcal{D}(\mathbb{P}_{G(s);A}(\cdot))$. Putting these two facts together, it follows that achieving (1) implies that outputs of G can be distinguished from uniformly random.

Formally, let us assume that there is an adversary \mathcal{D} that violates our hypothesis, i.e., that distinguishes some pair of inputs $(\mathbb{P}_{G(s);A}(\rho_{AB}), \mathbb{P}_{G(s);A}(\sigma_A \otimes \rho_B))$ with probability at least $1/p(n)$ for some polynomial p . Then we'll show an algorithm \mathcal{D}' , that breaks the pseudorandom generator G . On input $y \in \{0, 1\}^m$, algorithm \mathcal{D}' does the following:

- with probability $1/2$, run \mathcal{D} on input $\mathbb{P}_{y;A}(\rho_{AB})$;
- with probability $1/2$, run \mathcal{D} on input $\mathbb{P}_{y;A}(\sigma_A \otimes \rho_B)$.

Now if \mathcal{D} is able to correctly determine which of the cases we gave it, \mathcal{D}' decides that y must have been distributed pseudorandomly and outputs 1, else it decides that y is uniformly distributed and outputs 0.

Notice that if $y = G(s)$, by definition \mathcal{D}' outputs 1 when \mathcal{D} correctly distinguishes the two inputs, which occurs with probability at least $1/2 + 1/p(n)$ by the assumption on \mathcal{D} . On the other hand, suppose $y \xleftarrow{\$} \{0, 1\}^m$; then the register A is mapped to the maximally mixed state, and hence $\mathbb{P}_{y;A}(\rho_{AB}) = \mathbb{P}_{y;A}(\sigma_A \otimes \rho_B) = \mathbb{1}_A \otimes \rho_B$. In that case, \mathcal{D} is correct with probability at most $1/2 + \text{negl}(n)$ (indeed, this is true for any QPT.) We conclude that \mathcal{D}' distinguishes the case $y = G(s)$ from the case $y \xleftarrow{\$} \{0, 1\}^m$ with non-negligible probability; this contradicts the assumption that G is a qPRG. \square

Finally, to prove that the construction in [Scheme 2](#) is IND-CPA-secure, and thus establish [Theorem 3](#), it remains to extend the above proof to a slightly more general scenario. Recall that $\text{Enc}_i(\rho) = |f_i^{2n}(d)\rangle\langle f_i^{2n}(d)| \otimes P_r \rho P_r$ where $r = G(d)$. [Lemma 20](#) already shows that essentially no QPT adversary can distinguish $(P_r \otimes \mathbb{1}_E) \rho_{ME} (P_r \otimes \mathbb{1}_E)$ from $(P_r \otimes \mathbb{1}_E)(|0\rangle\langle 0| \otimes \rho_E)(P_r \otimes \mathbb{1}_E)$, for any efficiently preparable bipartite state ρ_{ME} over the message space and the environment. It remains to show that this indistinguishability still holds if the adversary is also provided the classical advice $f_i^{2n}(d)$. We can prove this extended indistinguishability by extending the hybrid argument in the proof of [Lemma 19](#) in a standard way. To sketch the argument, first recall that the ‘‘predictor’’ algorithm succeeds at predicting the $i + 1^{\text{st}}$ bit of $G(U_n)$ given as input the first i bits of the output of G . Now we also allow the predictor to read the bits of $f_i^{2n}(d)$. Success implies breaking the hard-core of f (which is used to define and

ensure the security of the qPRG G). We conclude that the states

$$|f_i^{2n}(d)\rangle\langle f_i^{2n}(d)| \otimes \mathbb{P}_{G(s);M}(\rho_{ME}) \quad \text{and} \quad |f_i^{2n}(d)\rangle\langle f_i^{2n}(d)| \otimes \mathbb{P}_{r';M}(\rho_{ME})$$

are computationally indistinguishable for uniformly random s, r' . The right-hand side encryption above obviously satisfies IND-CPA, so we also have computational indistinguishability of

$$|f_i^{2n}(d)\rangle\langle f_i^{2n}(d)| \otimes \mathbb{P}_{r';M}(\rho_{ME}) \quad \text{and} \quad |f_i^{2n}(d)\rangle\langle f_i^{2n}(d)| \otimes \mathbb{P}_{r';M}(|0\rangle\langle 0|_M \otimes \rho_E).$$

By transitivity of computational indistinguishability, we conclude that

$$|f_i^{2n}(d)\rangle\langle f_i^{2n}(d)| \otimes \mathbb{P}_{G(s);M}(|0\rangle\langle 0|_M \otimes \rho_E) \quad \text{and} \quad |f_i^{2n}(d)\rangle\langle f_i^{2n}(d)| \otimes \mathbb{P}_{G(s);M}(\rho_{ME}),$$

which completes the proof of [Theorem 3](#).

6 Conclusion

We have defined semantic security for the encryption of quantum data and shown its equivalence with indistinguishability; these results are given in the uniform model for quantum computations, but as is standard classically (see Chapter 5 of Goldreich’s text [\[Gol04a\]](#)), these definitions can be adjusted to the case of “non-uniform” (but still polynomial-time) adversaries, whose messages need not be generated efficiently. While the proof is somewhat different, the equivalence of IND and SEM still hold in this case. The constructions of encryption schemes (IND-CCA1 symmetric-key and IND-CPA public-key) presented above carry over as well, except that we now require primitives (qPRFs and qTOWPs, respectively) which are secure against non-uniform adversaries.

6.1 Extensions and Future Work

We now briefly discuss some possible extensions of the above results. In most cases, these extensions are a matter of modifying our definitions and proofs in a fairly straightforward way. We leave the other cases as interesting open problems.

- Our definitions of IND-CPA, IND-CCA1 and SEM assume that all of the relevant messages are generated in polynomial time. In other words, our results assume “uniform” adversaries. As is standard classically (see Chapter 5 of Goldreich’s text [\[Gol04a\]](#)), these definitions can be adjusted to the case of “non-uniform” (but still polynomial-time) adversaries, whose messages need not be generated efficiently. While the proof is of course somewhat different, the equivalence of IND and SEM still hold in this case. The encryption schemes (IND-CCA1 symmetric-key and IND-CPA public-key) presented above carry over as well, except that we now require primitives (qPRFs and qTOWPs, respectively) which are secure against non-uniform adversaries.

- Our symmetric-key encryption scheme assumes that the decryption algorithm measures a portion of the input in order to recover a classical randomness string, prior to decrypting. One might find this requirement suspicious, e.g., if a perfect measurement device is too much to assume. This requirement can be removed, but we then need to assume that the relevant primitives (OWFs and qPRFs) are secure against superposition queries. This can also be achieved (see [Zha12]).
- One outstanding open problem is to define and construct schemes for CCA2 (adaptive chosen ciphertext attack) security in the case of the encryption of quantum states. Classically, CCA2 security is defined as CCA1, with the further property that the adversary is allowed to query the decryption oracle even *after* the challenge query, *provided* he does not query about the challenge ciphertext itself (otherwise the challenger aborts the game.) The obvious way to define this in the quantum world is to require that every decryption query performed by the adversary after the challenge query is ‘very different’ from the challenge query itself (e.g., it is orthogonal to the challenge ciphertext.) But the problem here is that this condition might be impossible for the challenger to check: for example, the adversary might embed in a decryption query a component non-orthogonal to the challenge query, but with such a small amplitude that the challenger cannot detect it with high probability. Even if it is unclear whether this issue could raise problems in any actual reduction, it would be anyway a striking asymmetry to the classical case, because there would be no way for the challenger to check that the adversary actually fulfilled the required condition. Hence, giving a satisfactory definition for CCA2 security in the quantum world remains an interesting open problem.

6.2 Acknowledgements

G. A. was supported by a Sapere Aude grant of the Danish Council for Independent Research, the ERC Starting Grant “QMULT” and the CHIST-ERA project “CQC”. A. B. was supported by Canada’s NSERC. B.F. was supported by the Department of Defense. T. G. was supported by the German Federal Ministry of Education and Research (BMBF) within EC-SPRIDE and CROSSING. C. S. was supported by a 7th framework EU SIQS and a NWO VIDI grant. M. S.J was supported by the Ontario Ontario Graduate Scholarship Program. T. G. and C.S. would like to thank COST Action IC1306 for networking support. A. B., G. A., T. G., and C. S. would like to thank the organizers of the Dagstuhl Seminar 15371 “Quantum Cryptanalysis” for providing networking and useful interactions and support during the writing of this paper.

References

- Aar09. Scott Aaronson. Quantum copy-protection and quantum money. In *Computational Complexity, 2009. CCC’09. 24th Annual IEEE Conference on*, pages 229–242. IEEE, 2009.

- ABB⁺14. Romain Alléaume, Cyril Branciard, Jan Bouda, Thierry Debuisschert, Mehrdad Dianati, Nicolas Gisin, Mark Godfrey, Philippe Grangier, Thomas Länger, Norbert Lütkenhaus, Christian Monyk, Philippe Painchault, Momtchil Peev, Andreas Poppe, Thomas Pornin, John Rarity, Renato Renner, Gregoire Ribordy, Michel Riguidel, Louis Salvail, Andrew Shields, Harald Weinfurter, and Anton Zeilinger. Using quantum key distribution for cryptographic purposes: A survey. *Theoretical Computer Science*, 560:62–81, 2014.
- AC02. Mark Adcock and Richard Cleve. A quantum goldreich-levin theorem with cryptographic applications. In Helmut Alt and Afonso Ferreira, editors, *STACS 2002*, volume 2285 of *Lecture Notes in Computer Science*, pages 323–334. Springer Berlin Heidelberg, 2002.
- AC12. Scott Aaronson and Paul Christiano. Quantum money from hidden subspaces. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 41–60. ACM, 2012.
- AKN98. Dorit Aharonov, Alexei Kitaev, and Noam Nisan. Quantum circuits with mixed states. In *Proceedings of the thirtieth annual ACM symposium on Theory of computing*, pages 20–30. ACM, 1998.
- AMTdW00. Andris Ambainis, Michele Mosca, Alain Tapp, and Ronald de Wolf. Private quantum channels. In *Foundations of Computer Science, 2000. Proceedings. 41st Annual Symposium on*, pages 547–553, 2000.
- BB84. Charles Bennett and Gilles Brassard. Quantum cryptography: Public key distribution and coin tossing. In *Proceedings of the International Conference on Computers, Systems, and Signal Processing*, pages 175–179, 1984.
- BBD09. Daniel J Bernstein, Johannes Buchmann, and Erik Dahmen, editors. *Post-Quantum Cryptography*. Springer, 2009.
- BDF⁺11. Dan Boneh, Özgür Dagdelen, Marc Fischlin, Anja Lehmann, Christian Schaffner, and Mark Zhandry. Random oracles in a quantum world. In *Advances in Cryptology—ASIACRYPT 2011*, pages 41–69, 2011.
- BFK09. Anne Broadbent, Joseph Fitzsimons, and Elham Kashefi. Universal blind quantum computation. In *Foundations of Computer Science, 2009. FOCS'09. 50th Annual IEEE Symposium on*, pages 517–526. IEEE, 2009.
- BGS13. Anne Broadbent, Gus Gutoski, and Douglas Stebila. Quantum one-time programs. In *Advances in Cryptology—CRYPTO 2013*, pages 344–360. Springer, 2013.
- BJ15. Anne Broadbent and Stacey Jeffery. Quantum homomorphic encryption for circuits of low T -gate complexity. In *Crypto 2015*, pages 609–629, 2015.
- BOCG⁺06. Michael Ben-Or, Claude Crépeau, Daniel Gottesman, Avinatan Hassidim, and Adam Smith. Secure multiparty quantum computation with (only) a strict honest majority. In *Foundations of Computer Science, 2006. FOCS'06. 47th Annual IEEE Symposium on*, pages 249–260. IEEE, 2006.
- BR03. P. Oscar Boykin and Vwani Roychowdhury. Optimal encryption of quantum bits. *Physical review A*, 67(4):042317, 2003.
- Bro15. Anne Broadbent. Delegating private quantum computations. *Canadian Journal of Physics*, pages 941–946, Jun 2015.
- BS16. Anne Broadbent and Christian Schaffner. Quantum cryptography beyond quantum key distribution. *Designs, Codes and Cryptography*, 78:351–382, 2016.

- BZ13. Dan Boneh and Mark Zhandry. Secure Signatures and Chosen Ciphertext Security in a Quantum Computing World. In Ran Canetti and Juan A. Garay, editors, *Crypto 2013*, volume 8043 of *LNCS*, pages 361–379. Springer, 2013.
- DD10. Simon Pierre Desrosiers and Frédéric Dupuis. Quantum entropic security and approximate quantum encryption. *IEEE Transactions on Information Theory*, 56(7):3455–3464, 2010.
- Des09. Simon Pierre Desrosiers. Entropic security in quantum cryptography. *Quantum Information Processing*, 8(4):331–345, August 2009.
- DH76. Whitfield Diffie and Martin E. Hellman. New directions in cryptography. *IEEE transactions on Information Theory*, 22(6):644–654, 1976.
- DNS10. Frédéric Dupuis, Jesper Buus Nielsen, and Louis Salvail. Secure two-party quantum evaluation of unitaries against specious adversaries. In *Advances in Cryptology—CRYPTO 2010*, pages 685–706. Springer, 2010.
- DNS12. Frédéric Dupuis, Jesper Buus Nielsen, and Louis Salvail. Actively secure two-party evaluation of any quantum operation. In *Advances in Cryptology—CRYPTO 2012*, pages 794–811. Springer, 2012.
- FKS⁺13. Serge Fehr, Jonathan Katz, Fang Song, Hong-Sheng Zhou, and Vassilis Zikas. Feasibility and completeness of cryptographic tasks in the quantum world. In *Theory of Cryptography*, pages 281–296. Springer, 2013.
- GGM86. Oded Goldreich, Shafi Goldwasser, and Silvio Micali. How to construct random functions. *Journal of the ACM*, 33(4):792–807, 1986.
- GHS15. Tommaso Gagliardoni, Andreas Hülsing, and Christian Schaffner. Semantic security and indistinguishability in the quantum world, 2015. <http://arxiv.org/abs/1504.05255>.
- GL89. O. Goldreich and L. A. Levin. A hard-core predicate for all one-way functions. In *Proceedings of the Twenty-first Annual ACM Symposium on Theory of Computing*, STOC '89, pages 25–32, New York, NY, USA, 1989. ACM.
- GM84. Shafi Goldwasser and Silvio Micali. Probabilistic encryption. *Journal of Computer and System Sciences*, 28(2):270 – 299, 1984.
- Gol04a. Oded Goldreich. *Foundations of Cryptography: Volume 2, Basic Applications*. Cambridge University Press, Cambridge, UK, 2004.
- Gol04b. Oded Goldreich. *Foundations of Cryptography: Volume 2, Basic Applications*. Cambridge University Press, New York, NY, USA, 2004.
- GPV08. Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing*, STOC '08, pages 197–206, New York, NY, USA, 2008. ACM.
- HILL99. Johan Håstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. A pseudorandom generator from any one-way function. *SIAM J. Comput.*, 28:1364–1396, March 1999.
- HLSW04. Patrick Hayden, Debbie Leung, Peter W Shor, and Andreas Winter. Randomizing quantum states: Constructions and applications. *Communications in Mathematical Physics*, 250(2):371–391, 2004.
- KK07. Elham Kashefi and Iordanis Kerenidis. Statistical zero knowledge and quantum one-way functions. *Theoretical Computer Science*, 378(1):101 – 116, 2007.
- Kos07. Takeshi Koshihba. Security notions for quantum public-key cryptography. *IEICE TRANSACTIONS on Fundamentals of Electronics, Communications and Computer Sciences*, J90-A(5):367–375, Feb 2007.

- Leu02. Debbie W. Leung. Quantum vernam cipher. *Quantum Information and Computation*, 2(1):14–34, 2002.
- MRV07. C. Moore, A. Russell, and U. Vazirani. A classical one-way function to confound quantum adversaries. *eprint arXiv:quant-ph/0701115*, January 2007.
- MS10. Michele Mosca and Douglas Stebila. Quantum coins. *Error-Correcting Codes, Finite Geometries and Cryptography*, 523:35–47, 2010.
- OTU00. Tatsuaki Okamoto, Keisuke Tanaka, and Shigenori Uchiyama. Quantum public-key cryptosystems. In Mihir Bellare, editor, *Advances in Cryptology CRYPTO 2000*, volume 1880 of *Lecture Notes in Computer Science*, pages 147–165. Springer Berlin Heidelberg, 2000.
- PW08. Chris Peikert and Brent Waters. Lossy trapdoor functions and their applications. In *Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing, STOC '08*, pages 187–196, New York, NY, USA, 2008. ACM.
- Sha49. C. E. Shannon. Communication theory of secrecy systems*. *Bell System Technical Journal*, 28(4):656–715, Oct 1949.
- Sho94. Peter W. Shor. Algorithms for Quantum Computation: Discrete Logarithms and Factoring. In *FOCS 1994*, pages 124–134. IEEE Computer Society Press, 1994.
- Son14. Fang Song. A note on quantum security for post-quantum cryptography. In *Post-Quantum Cryptography*, pages 246–265. Springer, 2014.
- Unr10. Dominique Unruh. Universally composable quantum multi-party computation. In *Advances in Cryptology–EUROCRYPT 2010*, pages 486–505. Springer, 2010.
- Unr14. Dominique Unruh. Revocable quantum timed-release encryption. In *Advances in Cryptology–EUROCRYPT 2014*, pages 129–146. Springer, 2014.
- Unr15. Dominique Unruh. Non-interactive zero-knowledge proofs in the quantum random oracle model. In *Advances in Cryptology–EUROCRYPT 2015*, pages 755–784, 2015.
- Vel13. Maria Velema. Classical encryption and authentication under quantum attacks. Master’s thesis, Master of Logic, University of Amsterdam, 2013. <http://arxiv.org/abs/1307.3753>.
- Wie83. Stephen Wiesner. Conjugate coding. *ACM Sigact News*, 15(1):78–88, 1983.
- WZ82. W. K. Wootters and W. H. Zurek. A single quantum cannot be cloned. *Nature*, 299(5886):802–803, Oct 1982.
- XY12. Chong Xiang and Li Yang. Indistinguishability and semantic security for quantum encryption scheme. *Proc. SPIE*, 8554:85540G–85540G–8, 2012.
- Zha12. Mark Zhandry. How to Construct Quantum Random Functions. In *FOCS 2012*, pages 679–687. IEEE, 2012.

A Alternative Definitions of Quantum Security

Here, we present further definitions of quantum semantic security and indistinguishability, and prove their equivalence. The full chain of equivalencies is given in Figure 3.




Fig. 3. Relationship between security definitions.

A.1 SEM2

Definition 21 (Message-classical function generator). A message-classical function generator \mathcal{M} is a QPT message generator (as in SEM (Definition 8)) such that for each $pk \in \mathcal{K}_{pub}$ and ρ in the range of $\mathcal{M}(pk)$, there is some binary string y such that $|y\rangle \in \mathcal{H}_F$ and $\rho_{MEF} = \rho_{ME} \otimes |y\rangle \langle y|$.

That is, the F system is classical, unentangled from and uncorrelated with the rest of ρ .

In particular, $\rho_F = |y\rangle \langle y|$.

Definition 22 (SEM2). A qPKE scheme $(\text{KeyGen}, \text{Enc}, \text{Dec})$ is SEM2-secure if for any QPT adversary \mathcal{A} , there exists a QPT simulator \mathcal{S} such that for all message-classical function generators \mathcal{M} ,

$$\left| \Pr \left[\mathcal{A} \{ \text{Enc}_{pk} \otimes \mathbf{1}_E \} \rho_{ME} \right] = \rho_F \right] - \Pr \left[\mathcal{S}(\rho_E) = \rho_F \right] \right| \leq \text{negl}(n)$$

where the outputs of \mathcal{A} and \mathcal{S} are measured in the computational basis before equality is checked, $\rho_{MEF} \leftarrow \mathcal{M}(pk)$, and the probabilities are taken over $(pk, sk) \leftarrow \text{KeyGen}(1^n)$ and the internal randomness of Enc , \mathcal{A} , \mathcal{S} and \mathcal{D} .

- **SEM2-CPA:** In addition to the above, all QPTs are given oracle access to Enc_{pk} .
- **SEM2-CCA1:** In addition to SEM2-CPA, \mathcal{M} is given oracle access to Dec_{sk} .

Proposition 23. If a quantum encryption scheme $(\text{KeyGen}, \text{Enc}, \text{Dec})$ is semantically secure, then it is SEM2 secure.

Proof. A message-classical function generator is also a message. In SEM, have the distinguisher \mathcal{D} implement an equality test (simulate any efficient classical circuit implementing it; if the input lengths aren't the same, output 0 immediately). \square

A.2 SEM3

Definition 24 (Message Generator-Function Pair). A message generator-function pair is a tuple (\mathcal{M}, f) , such that \mathcal{M} is a QPT message generator (as in IND (Definition 7)) and $f = (f_n)_n$ is a QPT algorithm, such that $f_{pk} := f_n(pk)$ is the description of a boolean circuit, for $pk \in \mathcal{K}_{pub}$, with the number of input bits to f_{pk} equal to the number of measurement gates in the quantum circuit \mathcal{M}_n . In the symmetric-key scenario, f_n has no input.

Definition 25 (SEM3). A qPKE scheme $(\text{KeyGen}, \text{Enc}, \text{Dec})$ is SEM3-secure if for any QPT adversary \mathcal{A} , there exists a QPT simulator \mathcal{S} such that for all message generator-function pairs (\mathcal{M}, f) ,

$$\left| \Pr \left[\mathcal{A} \{ (\text{Enc}_{pk} \otimes \mathbf{1}_E) \rho_{ME} \} = f_{pk}(x) \right] - \Pr \left[\mathcal{S}(\rho_E) = f_{pk}(x) \right] \right| \leq \text{negl}(n)$$

where the outputs of \mathcal{A} and \mathcal{S} are measured in the computational basis before equality is checked, $\rho_{ME} \leftarrow \mathcal{M}(pk)$, x is the string of measurement results generating ρ_{ME} , and the probabilities are taken over $(pk, sk) \leftarrow \text{KeyGen}(1^n)$ and the internal randomness of Enc , \mathcal{A} , \mathcal{S} and \mathcal{D} .

- **SEM3-CPA:** In addition to the above, all QPTs are given oracle access to Enc_{pk} .
- **SEM3-CCA1:** In addition to SEM3-CPA, \mathcal{M} is given oracle access to Dec_{sk} .

We note that f_{pk} is a function of the random input and measurement results, which completely determine the state. Hence, if f_{pk} is the identity for all pk , and it can be computed given a ciphertext, this means we can compute measurement results necessary to prepare the state. Simulating the message generator but selecting for the correct measurement results would allow the preparation of the same state again, although this is not in general efficient.

Proposition 26 (SEM2 implies SEM3). If a quantum encryption scheme $(\text{KeyGen}, \text{Enc}, \text{Dec})$ is SEM2 secure, then it is SEM3 secure.

Proof. A message generator-function pair (\mathcal{M}, f) can be turned into a message-classical function generator \mathcal{M} by copying the measurement results x_1, x_2, \dots, x_m after each measurement gate, and applying f_{pk} to $x_1 x_2 \dots x_m$ and letting the result be the F system. \square

A.3 IND'

Definition 27 (IND'). A $qPKE$ scheme $(\text{KeyGen}, \text{Enc}, \text{Dec})$ is IND' secure if for every QPT adversary $\mathcal{A} = (\mathcal{M}, \mathcal{D})$ we have:

$$\Pr [\mathcal{D}\{(\text{Enc}_{pk} \otimes \mathbb{1}_E)\rho_{ME}^{(b)}\} = b] \leq \frac{1}{2} + \text{negl}(n)$$

where $\rho_{ME} \leftarrow \mathcal{M}(pk)$, for b a uniformly random bit, $\rho_{ME}^{(1)} = \rho_{ME}$ and $\rho_{ME}^{(0)} = |0\rangle\langle 0|_M \otimes \rho_E$, and the probabilities are taken over $(pk, sk) \leftarrow \text{KeyGen}(1^n)$, b and the internal randomness of Enc , \mathcal{M} , and \mathcal{D} .

- **IND'-CPA:** In addition to the above, \mathcal{M} and \mathcal{D} are given oracle access to Enc_{pk} .
- **IND'-CCA1:** In addition to IND'-CPA, \mathcal{M} is given oracle access to Dec_{sk} .

Proposition 28 (IND' \iff IND). A quantum encryption scheme $(\text{KeyGen}, \text{Enc}, \text{Dec})$ is IND' secure if and only if it is IND secure.

Proof. We drop brackets and the register subscripts where possible.

$$\begin{aligned} & \Pr[\mathcal{D}(\text{Enc}_{pk} \otimes \mathbb{1}_E)\rho^{(b)} = b] \\ &= \Pr[\mathcal{D}(\text{Enc}_{pk} \otimes \mathbb{1}_E)\rho^{(b)} = b \mid b = 1] \Pr[b = 1] \\ & \quad + \Pr[\mathcal{D}(\text{Enc}_{pk} \otimes \mathbb{1}_E)\rho^{(b)} = b \mid b = 0] \Pr[b = 0] \\ &= \frac{1}{2}(\Pr[\mathcal{D}(\text{Enc}_{pk} \otimes \mathbb{1}_E)\rho = 1] + \Pr[\mathcal{D}(\text{Enc}_{pk} |0\rangle\langle 0| \otimes \rho_E) = 0]) \\ &\leq \frac{1}{2}(\Pr[\mathcal{D}(\text{Enc}_{pk} \otimes \mathbb{1}_E)\rho = 1] + 1 - \Pr[\mathcal{D}(\text{Enc}_{pk} |0\rangle\langle 0| \otimes \rho_E) = 1]) \\ &= \frac{1}{2} + \frac{1}{2}(\Pr[\mathcal{D}(\text{Enc}_{pk} \otimes \mathbb{1}_E)\rho = 1] - \Pr[\mathcal{D}(\text{Enc}_{pk} |0\rangle\langle 0| \otimes \rho_E) = 1]) \end{aligned}$$

Note that we only get \leq since \mathcal{D} may output some binary string other than 0 or 1. So:

$$\begin{aligned} & \Pr[\mathcal{D}(\text{Enc}_{pk} \otimes \mathbb{1}_E)\rho^{(b)} = b] - \frac{1}{2} \\ & \leq \frac{1}{2} |\Pr[\mathcal{D}((\text{Enc}_{pk} \otimes \mathbb{1}_E)\rho) = 1] - \Pr[\mathcal{D}(\text{Enc}_{pk} |0\rangle\langle 0| \otimes \rho_E) = 1]|. \end{aligned}$$

Thus, $\text{IND} \implies \text{IND}'$.

Now consider replacing \mathcal{D} with the distinguisher which starts the same as \mathcal{D} , but if \mathcal{D} would have output something other than 0 or 1, it simply outputs 0. Then the quantity $|\Pr[\mathcal{D}((\text{Enc}_{pk} \otimes \mathbb{1}_E)(\rho)) = 1] - \Pr[\mathcal{D}(\text{Enc}_{pk} |0\rangle\langle 0| \otimes \rho_E) = 1]|$ is the same for this new distinguisher, so without loss of generality, \mathcal{D} only outputs 0 or 1.

Then the first \leq becomes an $=$, i.e.

$$\begin{aligned} & \Pr[\mathcal{D}(\text{Enc}_{pk} \otimes \mathbb{1}_E)\rho^{(b)} = b] - \frac{1}{2} \\ &= \frac{1}{2}(\Pr[\mathcal{D}(\text{Enc}_{pk} \otimes \mathbb{1}_E)\rho = 1] - \Pr[\mathcal{D}(\text{Enc}_{pk} |0\rangle\langle 0| \otimes \rho_E) = 1]) \end{aligned}$$

and, similarly,

$$\begin{aligned}
& \Pr[\mathcal{D}(\text{Enc}_{pk} \otimes \mathbb{1}_E)\rho^{(b)} = b \oplus 1] \\
&= \Pr[\mathcal{D}(\text{Enc}_{pk} \otimes \mathbb{1}_E)\rho^{(b)} = b \oplus 1 \mid b = 1] \Pr[b = 1] \\
&\quad + \Pr[\mathcal{D}(\text{Enc}_{pk} \otimes \mathbb{1}_E)\rho^{(b)} = b \oplus 1 \mid b = 0] \Pr[b = 0] \\
&= \frac{1}{2}(\Pr[\mathcal{D}(\text{Enc}_{pk} \otimes \mathbb{1}_E)\rho = 0] + \Pr[\mathcal{D}(\text{Enc}_{pk} | 0\rangle \langle 0| \otimes \rho_E) = 1]) \\
&= \frac{1}{2}(1 - \Pr[\mathcal{D}(\text{Enc}_{pk} \otimes \mathbb{1}_E)\rho = 1] + \Pr[\mathcal{D}(\text{Enc}_{pk} | 0\rangle \langle 0| \otimes \rho_E) = 1]) \\
&= \frac{1}{2} + \frac{1}{2}(\Pr[\mathcal{D}(\text{Enc}_{pk} | 0\rangle \langle 0| \otimes \rho_E) = 1] - \Pr[\mathcal{D}(\text{Enc}_{pk} \otimes \mathbb{1}_E)\rho = 1])
\end{aligned}$$

so,

$$\begin{aligned}
& \Pr[\mathcal{D}(\text{Enc}_{pk} \otimes \mathbb{1}_E)\rho^{(b)} = b \oplus 1] - \frac{1}{2} \\
&= \frac{1}{2}(\Pr[\mathcal{D}(\text{Enc}_{pk} | 0\rangle \langle 0| \otimes \rho_E) = 1] - \Pr[\mathcal{D}(\text{Enc}_{pk} \otimes \mathbb{1}_E)\rho = 1])
\end{aligned}$$

Combining the above,

$$\begin{aligned}
& \frac{1}{2} |\Pr[\mathcal{D}(\text{Enc}_{pk} \otimes \mathbb{1}_E)\rho = 1] - \Pr[\mathcal{D}(\text{Enc}_{pk} | 0\rangle \langle 0| \otimes \rho_E) = 1]| \\
&= \max\{\Pr[\mathcal{D}(\text{Enc}_{pk} \otimes \mathbb{1}_E)\rho^{(b)} = b] - \frac{1}{2}, \Pr[\mathcal{D}(\text{Enc}_{pk} \otimes \mathbb{1}_E)\rho^{(b)} = b \oplus 1] - \frac{1}{2}\}
\end{aligned}$$

Hence $\text{IND}' \implies \text{IND}$ by applying IND to both \mathcal{D} and $\mathcal{D} \oplus 1$ (the latter outputs the answer opposite to \mathcal{D}), for $\Pr[\mathcal{D}(\text{Enc}_{pk} \otimes \mathbb{1}_E)\rho^{(b)} = b]$ and $\Pr[\mathcal{D}(\text{Enc}_{pk} \otimes \mathbb{1}_E)\rho^{(b)} = b \oplus 1]$, respectively. The maximum of two negligible functions is again negligible. \square

Theorem 29 (SEM3 \implies IND'). *If a quantum encryption scheme (KeyGen, Enc, Dec) is SEM3 secure, then it is IND' secure.*

Proof. We drop brackets and the register subscripts where possible.

Let $(\mathcal{M}, \mathcal{D})$ be an IND' adversary.

Let us consider the SEM3 message generator \mathcal{M}' which runs $\rho_{ME} \leftarrow \mathcal{M}$ and outputs (with probability $\frac{1}{2}$ each) either the state ρ_{ME} or the state $|0\rangle \langle 0|_{\mathcal{M}} \otimes \rho_E$, and we denote its output by ρ'_{ME} . In particular, it prepares a random bit b to do so by measuring an ancillary qubit to which the Hadamard was applied.

Define $f_{pk}(xb) = b$.

Define the SEM3 adversary $\mathcal{A} := \mathcal{D}$.

In this way, the SEM3 game simulates the indistinguishability game, and

$$\Pr[\mathcal{A}((\text{Enc}_{pk} \otimes \mathbb{1}_E)(\rho'_{ME})) = f_{pk}(xb)] = \Pr[\mathcal{D}((\text{Enc}_{pk} \otimes \mathbb{1}_E)(\rho^{(b)})) = b]$$

Now, by SEM3, there is some simulator \mathcal{S} for \mathcal{A} so that

$$|\Pr[\mathcal{A}(\text{Enc}_{pk} \otimes \mathbb{1}_E)\rho'_{ME} = f_{pk}(xb)] - \Pr[\mathcal{S}\rho'_E = f_{pk}(xb)]| \leq \text{negl}(n)$$

i.e.

$$|\Pr[\mathcal{D}(\text{Enc}_{pk} \otimes \mathbf{1}_E)(\rho^{(b)}) = b] - \Pr[\mathcal{S}\rho'_E = f_{pk}(xb)]| \leq \text{negl}(n)$$

Note that \mathcal{S} 's input ρ'_E is independent of b . Hence

$$\Pr[\mathcal{S}\rho'_E = f_{pk}(xb)] \leq \frac{1}{2}.$$

Finally, by the triangle inequality applied to the last two inequalities,

$$\Pr[\mathcal{D}(\text{Enc}_{pk} \otimes \mathbf{1}_E)\rho^{(b)} = b] \leq \frac{1}{2} + \text{negl}(n)$$

□