Postselection Technique for Quantum Channels with Applications to Quantum Cryptography

Matthias Christandl
Faculty of Physics, Ludwig-Maximilians-Universität München, Theresienstrasse 37, 80333 Munich, Germany

Robert König
Institute for Quantum Information, California Institute of Technology, Pasadena, California 91125, USA

Renato Renner
Institute for Theoretical Physics, ETH Zurich, 8093 Zurich, Switzerland
(Received 17 September 2008; published 14 January 2009)

We propose a general method for studying properties of quantum channels acting on an n-partite system, whose action is invariant under permutations of the subsystems. Our main result is that, in order to prove that a certain property holds for an arbitrary input, it is sufficient to consider the case where the input is a particular de Finetti-type state, i.e., a state which consists of n identical and independent copies of an (unknown) state on a single subsystem. Our technique can be applied to the analysis of information-theoretic problems. For example, in quantum cryptography, we get a simple proof for the fact that security of a discrete-variable quantum key distribution protocol against collective attacks implies security of the protocol against the most general attacks. The resulting security bounds are tighter than previously known bounds obtained with help of the exponential de Finetti theorem.

DOI: 10.1103/PhysRevLett.102.020504 PACS numbers: 03.67.Hk, 02.20.Qs, 03.67.Dd

In quantum mechanics, the most general way of describing the evolution of a subsystem A (A may be part of a larger system) at time t to a subsystem A' at a later point in time t' is by application of a quantum channel. Mathematically, a quantum channel is a completely positive trace-preserving (CPTP) map transforming the reduced density matrix ρ_A of system A at time t to $\rho_{A'}$, the reduced density matrix of system A' at time t'. CPTP maps are used in various areas of physics and information theory. A CPTP map modeling a particular quantum communication channel, for instance, describes how the channel output $\rho_{A'}$ depends on the input ρ_A.

A common method to characterize a given CPTP map $\mathcal{E}$ is to compare it to an idealized CPTP map $\mathcal{F}$ that is well understood, e.g., because it has a simple description. For instance, given a physical communication channel specified by $\mathcal{E}$, one may characterize its ability to reliably transmit messages by showing its similarity to a perfect channel $\mathcal{F}$ characterized by the identity mapping id. Another example is the analysis of information-theoretic or cryptographic protocols (e.g., for quantum key distribution). Here, $\mathcal{E}$ may be the action of the actual protocol while $\mathcal{F}$ is the ideal functionality the protocol is supposed to reproduce. We are then typically interested in proving that $\mathcal{E}$ is almost equal to $\mathcal{F}$ (in quantum cryptography, this corresponds to proving security).

In order to compare two CPTP maps $\mathcal{E}$ and $\mathcal{F}$, we need a notion of distance. A natural choice is the metric induced by the diamond norm $\|\cdot\|_\diamond$ [1] since it is directly related to the maximum probability that a difference can be observed between the processes described by $\mathcal{E}$ and $\mathcal{F}$, respectively. More precisely, consider a hypothetical game where a player is asked to guess whether a given physical process is described by $\mathcal{E}$ or $\mathcal{F}$, which are both equally likely to be the correct descriptions. If the player is allowed to observe the process once (with an input of his choice, possibly correlated with a reference system) then the maximum probability p of a correct guess is given by $p = \frac{1}{2} + \frac{1}{4}\|\mathcal{E} - \mathcal{F}\|_\diamond$. In particular, if $\mathcal{E}$ and $\mathcal{F}$ are identical, the distance equals zero and, hence, $p = \frac{1}{2}$, corresponding to a random guess. On the other hand, if $\mathcal{E}$ and $\mathcal{F}$ are perfectly distinguishable, we have $\|\mathcal{E} - \mathcal{F}\|_\diamond = 2$ and $p = 1$.

Here, we present a general method for computing an upper bound on the distance $\|\mathcal{E} - \mathcal{F}\|_\diamond$ between two maps $\mathcal{E}$ and $\mathcal{F}$, provided they act symmetrically on an n-partite system with subsystems $\mathcal{H}$ of finite dimension. While, by definition, the diamond norm involves a maximization over all possible inputs, we show that for calculating the bound it is sufficient to consider (relative to a reference system) the particular input

$$\tau_{\mathcal{H}^n} = \int \sigma_{\mathcal{H}}^{\otimes n} \mu(\sigma_{\mathcal{H}}), \tag{1}$$

where $\mu(\cdot)$ is the measure on the space of density operators on a single subsystem induced by the Hilbert-Schmidt metric. States of the form (1) are also known as de Finetti states. They describe the joint state of n subsystems prepared as identical and independent copies of an (unknown) density operator $\sigma_{\mathcal{H}}$. Because of their structure, de Finetti states are usually easy to handle in calculations and proofs, as outlined below.

As an example, we apply this result to the security analysis of quantum key distribution (QKD) schemes

[3,4]. Let $\mathcal{E}$ be the map describing a given QKD protocol, which takes as input n predistributed particle pairs (which may have been generated in a preliminary protocol step). Security of the protocol (against the most general attacks) is then defined by the requirement that the protocol $\mathcal{E}$ is close to the ideal functionality $\mathcal{F}$ that simply outputs a perfect key, independently of the input (which may be arbitrarily compromised by the action of an adversary). Now, according to our main result, this distance is bounded by simply evaluating the map $\mathcal{E}$ for an input of the form (1) and comparing the generated key with a perfect key. We further show that the latter is equivalent to proving security of the scheme against a restricted type of attacks, called collective attacks, where the adversary is assumed to attack each of the particle pairs independently and identically. Our result thus gives a simple proof for the statement (proved originally in [5,6]) that security of a QKD protocol against collective attacks implies security against the most general attacks. The resulting security bounds are tighter than previously known bounds obtained by proofs relying on the exponential de Finetti theorem [5].

Main result.—Let Δ be a linear map from $\mathrm{End}(\mathcal{H}^{\otimes n})$ to $\mathrm{End}(\mathcal{H}')$. In particular, Δ may be the difference between two CPTP maps. $\mathrm{End}(\mathcal{L})$ denotes the space of all endomorphisms on $\mathcal{L}$, which includes the density operators on $\mathcal{L}$. We denote by π the map on $\mathrm{End}(\mathcal{H}^{\otimes n})$ that permutes the subsystems with permutation π [7]. Our main result, the post-selection theorem [8], gives an upper bound on the norm of a permutation-invariant map in terms of the action of the map on a purification [9] $\tau_{\mathcal{H}^n\mathcal{R}}$ of the state $\tau_{\mathcal{H}^n}$ defined by (1).

Theorem 1. If for any permutation π there exists a linear CPTP map $\mathcal{K}_\pi$ such that $\Delta \circ \pi = \mathcal{K}_\pi \circ \Delta$, then

$$\|\Delta\|_\diamond \leq g_{n,d} \|(\Delta \otimes \mathrm{id})(\tau_{\mathcal{H}^n\mathcal{R}})\|_1.$$

id denotes the identity map on $\mathrm{End}(\mathcal{R})$ and

$$g_{n,d} = \binom{n + d^2 - 1}{n} \leq (n+1)^{d^2-1},$$

for $d = \dim\mathcal{H}$.

Compared to the left-hand side of this bound, which involves an optimization over all input states and is generally hard to evaluate, the right-hand side is significantly easier to deal with since it only involves a single, very specific input state.

The proof of Theorem 1 uses the following lemma which relates arbitrary density operators $\rho_{\mathcal{H}^n\mathcal{K}^n}$ on the symmetric subspace $\mathrm{Sym}^n(\mathcal{H}\otimes\mathcal{K}) \subset (\mathcal{H}\otimes\mathcal{K})^{\otimes n}$, for $\mathcal{K} \cong \mathcal{H}$, to a particular purification of $\tau_{\mathcal{H}^n}$. We define the state $\tau_{\mathcal{H}^n\mathcal{K}^n} = \int \sigma_{\mathcal{H}\mathcal{K}}^{\otimes n} d(\sigma_{\mathcal{H}\mathcal{K}})$ on $\mathrm{Sym}^n(\mathcal{H}\otimes\mathcal{K})$, where $d(\cdot)$ is the measure on the pure states induced by the Haar measure on the unitary group acting on $\mathcal{H}\otimes\mathcal{K}$. We note that $\tau_{\mathcal{H}^n\mathcal{K}^n}$ extends the state $\tau_{\mathcal{H}^n}$ defined in (1), i.e., $\mathrm{tr}_{\mathcal{K}^n}\tau_{\mathcal{H}^n\mathcal{K}^n} = \tau_{\mathcal{H}^n}$; the measure $\mu(\cdot)$ furthermore is the one induced by the Hilbert-Schmidt metric on $\mathrm{End}(\mathcal{H})$ [10]. Let now $\tau_{\mathcal{H}^n\mathcal{K}^n\mathcal{N}}$ be a purification of $\tau_{\mathcal{H}^n\mathcal{K}^n}$.

Lemma 2. For $\rho_{\mathcal{H}^n\mathcal{K}^n}$ a density operator supported on $\mathrm{Sym}^n(\mathcal{H}\otimes\mathcal{K})$, with $\mathcal{K}\cong\mathcal{H}$, there exists a linear completely positive trace-nonincreasing map $\mathcal{T}$ from the purifying system $\mathrm{End}(\mathcal{N})$ to $\mathbb{C}$ such that

$$\rho_{\mathcal{H}^n\mathcal{K}^n} = g_{n,d}(\mathrm{id}\otimes\mathcal{T})(\tau_{\mathcal{H}^n\mathcal{K}^n\mathcal{N}}), \tag{2}$$

where id is the identity map on $\mathrm{End}((\mathcal{H}\otimes\mathcal{K})^{\otimes n})$ and $d = \dim\mathcal{H}$.

Proof. Let $\mathcal{N} \cong \mathrm{Sym}^n(\mathcal{H}\otimes\mathcal{K})$ and let $\{|\nu_i\rangle\}_i$ be an eigenbasis of $\rho_{\mathcal{H}^n\mathcal{K}^n}$. Since, by Schur's lemma, $\tau_{\mathcal{H}^n\mathcal{K}^n}$ is the state proportional to the identity on $\mathrm{Sym}^n(\mathcal{H}\otimes\mathcal{K})$, $\tau_{\mathcal{H}^n\mathcal{K}^n\mathcal{N}} := |\Psi\rangle\langle\Psi|_{\mathcal{H}^n\mathcal{K}^n\mathcal{N}}$ is a purification of $\tau_{\mathcal{H}^n\mathcal{K}^n}$, where $|\Psi\rangle_{\mathcal{H}^n\mathcal{K}^n\mathcal{N}} := g_{n,d}^{-(1/2)}\sum_i |\nu_i\rangle\otimes|\nu_i\rangle$, and $g_{n,d}$ is the dimension of $\mathrm{Sym}^n(\mathcal{H}\otimes\mathcal{K})$. Furthermore, for any basis vector $|\nu_i\rangle$,

$$|\nu_i\rangle\langle\nu_i|_{\mathcal{H}^n\mathcal{K}^n} = g_{n,d}\,\mathrm{tr}_{\mathcal{N}}(\tau_{\mathcal{H}^n\mathcal{K}^n\mathcal{N}}\mathbb{1}_{\mathcal{H}^n\mathcal{K}^n}\otimes|\nu_i\rangle\langle\nu_i|_{\mathcal{N}}),$$

where $\mathbb{1}_{\mathcal{H}^n\mathcal{K}^n} \in \mathrm{End}((\mathcal{H}\otimes\mathcal{K})^{\otimes n})$ is the identity. This implies (2) with $\mathcal{T} : \sigma \mapsto \mathrm{tr}(\sigma\rho_{\mathcal{N}})$, since $\{|\nu_i\rangle\}_i$ is an eigenbasis of $\rho_{\mathcal{H}^n\mathcal{K}^n}$. Because $\mathcal{T}$ is clearly trace-nonincreasing, this concludes the proof. □

Proof of Theorem 1. We need to show that for any finite-dimensional space $\mathcal{R}'$ and any density operator $\rho_{\mathcal{H}^n\mathcal{R}'}$,

$$\|(\Delta\otimes\mathrm{id})(\rho_{\mathcal{H}^n\mathcal{R}'})\|_1 \leq g_{n,d}\|(\Delta\otimes\mathrm{id})(\tau_{\mathcal{H}^n\mathcal{R}})\|_1, \tag{3}$$

for some purification $\tau_{\mathcal{H}^n\mathcal{R}}$ of $\tau_{\mathcal{H}^n}$. In a first step, we show that it is sufficient to prove (3) for density operators $\rho_{\mathcal{H}^n\mathcal{R}'}$ with support on $\mathrm{Sym}^n(\mathcal{H}\otimes\mathcal{K})$, where $\mathcal{K}\cong\mathcal{H}$ and $\mathcal{R}' = \mathcal{K}^{\otimes n}$. To see this, let $\rho_{\mathcal{H}^n\mathcal{R}'}$ be an arbitrary density operator and define the density operator

$$\bar\rho_{\mathcal{H}^n\mathcal{R}'\mathcal{R}''} = \frac{1}{n!}\sum_\pi (\pi\otimes\mathrm{id})(\rho_{\mathcal{H}^n\mathcal{R}'})\otimes|\pi\rangle\langle\pi|_{\mathcal{R}''},$$

where the sum ranges over all permutations π of the n subsystems and where $\{|\pi\rangle\}_\pi$ is an orthonormal family of vectors on an auxiliary space $\mathcal{R}''$. Then, by construction, the reduced state $\bar\rho_{\mathcal{H}^n} = \mathrm{tr}_{\mathcal{R}'\mathcal{R}''}(\bar\rho_{\mathcal{H}^n\mathcal{R}'\mathcal{R}''})$ is permutation-invariant. Hence, according to [6,11], there exists a purification $\bar\rho_{\mathcal{H}^n\mathcal{H}^n}$ of $\bar\rho_{\mathcal{H}^n}$ supported on $\mathrm{Sym}^n(\mathcal{H}\otimes\mathcal{K})$. In particular, because all purifications are equivalent up to isometries, there exists a CPTP map $\mathcal{G}$ from $\mathrm{End}(\mathcal{K}^{\otimes n})$ to $\mathrm{End}(\mathcal{R}'\otimes\mathcal{R}'')$ such that $\bar\rho_{\mathcal{H}^n\mathcal{R}'\mathcal{R}''} = (\mathrm{id}\otimes\mathcal{G})(\bar\rho_{\mathcal{H}^n\mathcal{K}^n})$. Making use of the assumption on the permutation-invariance of Δ, we thus find that $\|(\Delta\otimes\mathrm{id})(\rho_{\mathcal{H}^n\mathcal{R}'})\|_1$ equals

$$\begin{aligned}\frac{1}{n!}\sum_\pi\|((\Delta\circ\pi)\otimes\mathrm{id})(\rho_{\mathcal{H}^n\mathcal{R}'})\|_1 &= \|(\Delta\otimes\mathrm{id})(\bar\rho_{\mathcal{H}^n\mathcal{R}'\mathcal{R}''})\|_1\\ &= \|(\Delta\otimes\mathcal{G})(\bar\rho_{\mathcal{H}^n\mathcal{K}^n})\|_1\\ &\leq \|(\Delta\otimes\mathrm{id})(\bar\rho_{\mathcal{H}^n\mathcal{K}^n})\|_1,\end{aligned}$$

where the last inequality holds because a CPTP map cannot

increase the norm. It thus remains to show that (3) holds for states $\rho_{\mathcal{H}^n\mathcal{K}^n}$ in $\mathrm{Sym}^n(\mathcal{H}\otimes\mathcal{K})$. By Lemma 2 there exists a map $\mathcal{T}$ such that $\rho_{\mathcal{H}^n\mathcal{K}^n} = g_{n,d}(\mathrm{id}\otimes\mathcal{T})(\tau_{\mathcal{H}^n\mathcal{K}^n\mathcal{N}})$. Then, by linearity, we have

$$\|(\Delta\otimes\mathrm{id})(\rho_{\mathcal{H}^n\mathcal{K}^n})\|_1 = g_{n,d}\|(\Delta\otimes\mathcal{T})(\tau_{\mathcal{H}^n\mathcal{K}^n\mathcal{N}})\|_1.$$

Inequality (3) then follows from the fact that $\mathcal{T}$ cannot increase the norm and by setting $\mathcal{R} = \mathcal{K}^{\otimes n}\otimes\mathcal{N}$. □

Application to quantum key distribution.—QKD is the art of generating a secret key known only to two distant parties, Alice and Bob, connected by an insecure quantum communication channel and an authentic classical channel [12]. Most QKD protocols can be subdivided into two parts. In the first, Alice and Bob use the quantum channel to distribute n entangled particle pairs (this phase may include advanced quantum protocols such as quantum repeaters). In the second part, they apply local measurements (we will restrict ourselves to the typical case of measurements that are independent and identical on each of the n pairs) followed by a sequence of classical postprocessing steps (such as parameter estimation, error correction, and privacy amplification) to extract ℓ key bits [13]. It induces a map $\mathcal{E}$ from $(\mathcal{H}_A\otimes\mathcal{H}_B)^{\otimes n}$ (the n particle pairs) to the set of pairs (S_A, S_B) of ℓ-bit strings (Alice and Bob's final keys, respectively) and C, where C is a transcript of the classical communication. Note that ℓ may depend on the input; in particular, $\ell = 0$ if the entanglement of the initial particle pairs is too small for key extraction.

A QKD protocol is said to be ε secure (for some small $\varepsilon \geq 0$) if, for any attack of an adversary, the final keys S_A and S_B computed by Alice and Bob are identical, uniformly distributed, and independent of the adversary's knowledge, except with probability ε. This criterion can be reformulated as a condition on the map $\mathcal{E}$. Since an adversary may have full control over the quantum channel connecting Alice and Bob, we require that, for any input to $\mathcal{E}$, the output is a pair (S_A, S_B) of secure keys of length $\ell \geq 0$ [15]. To make this more precise, let $\mathcal{S}$ be the map that acts on the output (S_A, S_B, C) of $\mathcal{E}$ by replacing (S_A, S_B) by a pair (S'_A, S'_B) of identical and uniformly distributed keys of the same length, while leaving C unchanged. With this definition, the concatenated map $\mathcal{F} := \mathcal{S}\circ\mathcal{E}$ describes an ideal key distillation scheme which always outputs a perfect key pair. We then say that $\mathcal{E}$ is ε secure if $\|\mathcal{E} - \mathcal{F}\|_\diamond \leq \varepsilon$.

$\mathcal{E}$ is typically invariant under permutations of the inputs. However, if it is not, permutation-invariance can be enforced by prepending an additional symmetrization step where both Alice and Bob permute their inputs according to a permutation $\bar{\pi}$ chosen at random by one party and communicated to the other using the classical channel [16]. We can thus apply Theorem 1 with $\Delta := \mathcal{E} - \mathcal{F}$, which implies that $\mathcal{E}$ is ε secure whenever

$$\|((\mathcal{E}-\mathcal{F})\otimes\mathrm{id})(\tau_{\mathcal{H}^n\mathcal{R}})\|_1 \leq \bar{\varepsilon} := \varepsilon(n+1)^{-(d^2-1)}, \quad (4)$$

where $\mathcal{H} := \mathcal{H}_A\otimes\mathcal{H}_B$, where $d = \dim(\mathcal{H})$, and where $\tau_{\mathcal{H}^n\mathcal{R}}$ is a purification of the state $\tau_{\mathcal{H}^n}$ defined by (1).

We will now employ (4) to show that for proving security of a QKD protocol it suffices to consider collective attacks, where the adversary acts on each of the signals independently and identically. Using the above formalism, we say that $\mathcal{E}$ is $\bar{\varepsilon}$ secure against collective attacks if $\|((\mathcal{E}-\mathcal{F})\otimes\mathrm{id})(\sigma^{\otimes n}_{\mathcal{H}\mathcal{K}})\|_1 \leq \bar{\varepsilon}$, for any (pure) $\sigma_{\mathcal{H}\mathcal{K}}$ on $\mathcal{H}\otimes\mathcal{K}$, where $\mathcal{K}\cong\mathcal{H}$. This immediately implies that the same bound holds for the extension $\tau_{\mathcal{H}^n\mathcal{K}^n} = \int\sigma^{\otimes n}_{\mathcal{H}\mathcal{K}}d(\sigma_{\mathcal{H}\mathcal{K}})$ of $\tau_{\mathcal{H}^n}$,

$$\begin{aligned}\|(\mathcal{E}-\mathcal{F})\otimes\mathrm{id}_{\mathcal{K}^n}(\tau_{\mathcal{H}^n\mathcal{K}^n})\|_1 &\leq \max_{\sigma_{\mathcal{H}\mathcal{K}}}\|(\mathcal{E}-\mathcal{F})\\ &\quad\otimes\mathrm{id}_{\mathcal{K}^n}(\sigma^{\otimes n}_{\mathcal{H}\mathcal{K}})\|_1\\ &\leq \bar{\varepsilon}. \quad (5)\end{aligned}$$

To obtain criterion (4), we need to show that a similar bound still holds if we consider a purification $\tau_{\mathcal{H}^n\mathcal{K}^n\mathcal{N}}$ of $\tau_{\mathcal{H}^n\mathcal{K}^n}$. For this, we think of $\mathcal{N}$ as an additional system that is available to an adversary. Because $\mathcal{N}$ can be chosen isomorphic to $\mathrm{Sym}^n(\mathcal{H}\otimes\mathcal{K})$, its dimension is bounded by $(n+1)^{d^2-1}$. The idea is then to compensate the extra information available to the adversary by slightly reducing the size of the final key. More precisely, according to the privacy amplification theorem (Theorem 5.5.1 of [6]), the protocol $\mathcal{E}'$ obtained from $\mathcal{E}$ by shortening the output of the hashing by $2\log_2\dim\mathcal{N} \leq 2(d^2-1)\log_2(n+1)$ bits satisfies

$$\begin{aligned}\|(\mathcal{E}'-\mathcal{F}')\otimes\mathrm{id}_{\mathcal{K}^n\mathcal{N}}(\kappa)\|_1 \leq \|(\mathcal{E}-\mathcal{F})\otimes\mathrm{id}_{\mathcal{K}^n}\\ \otimes\mathrm{tr}_{\mathcal{N}}(\kappa)\|_1.\end{aligned}$$

Setting κ equal to $\tau_{\mathcal{H}^n\mathcal{K}^n\mathcal{N}}$ and using (5), we conclude that $\|(\mathcal{E}'-\mathcal{F}')\otimes\mathrm{id}_{\mathcal{K}^n\mathcal{N}}(\tau_{\mathcal{H}^n\mathcal{K}^n\mathcal{N}})\|_1 \leq \bar{\varepsilon}$, which corresponds to (4). We have thus shown that $\bar{\varepsilon}$ security of $\mathcal{E}$ against collective attacks implies ε security of $\mathcal{E}'$ against general attacks.

In the security analysis against collective attacks, the security parameter $\bar{\varepsilon}$ can be chosen exponentially small, i.e., $\bar{\varepsilon} \leq 2^{-c\delta^2 n}$ (for some $c > 0$), at the only cost of reducing the key size by an (arbitrarily small) fraction δ compared to the asymptotically optimal rate. The crucial observation made in this Letter is that the security parameter ε for general attacks and $\bar{\varepsilon}$ are polynomially related [see (4)]. We thus find $\varepsilon \leq 2^{-c\delta^2 n+(d^2-1)\log_2(n+1)}$, which shows that security under the assumption of collective attacks implies full security essentially without changing the security parameter. This security estimate improves on previous estimates based on the exponential de Finetti theorem [17]. It can be applied to the security analysis of QKD in the practically relevant case where the capacity of the quantum channel is small.

Concluding remark.—The technical results in this Letter deal with quantum states and channels that commute with the action of the symmetric group on $\mathcal{H}^{\otimes n}$, but can be easily generalized to the action of an arbitrary finite or locally compact group G on a space $\mathcal{V}$ [18]. Seen in the light of more general symmetry groups G, we thus hope that our results will find fundamental applications in quantum physics beyond their presented use in quantum information theory.

R. K. acknowledges support from NSF Grants No. PHY-0456720 and No. PHY-0803371. R. R. received support from the EU project SECOQC.

[1] The diamond norm is given by $\|\mathcal{E}\|_\diamond = \sup_{k\in\mathbb{N}}\|\mathcal{E}\otimes \mathrm{id}_k\|_1$, where $\|\mathcal{F}\|_1 := \sup_{\|\sigma\|_1\leq 1}\|\mathcal{F}(\sigma)\|_1$, and $\|\sigma\|_1 := \mathrm{tr}\sqrt{\sigma^\dagger\sigma}$ is the trace norm. id_k denotes the identity map on states of a k-dimensional quantum system. The suprema are reached for positive σ and k equal to the dimension of the input of $\mathcal{E}$ [2].

[2] A. Y. Kitaev, Russ. Math. Surv. **52**, 1191 (1997).

[3] C. H. Bennett and G. Brassard, in *Proceedings of the IEEE International Conference on Computers, Systems and Signal Processing, 1984* (IEEE, Piscataway, NJ, 1984), pp. 175–179.

[4] A. K. Ekert, Phys. Rev. Lett. **67**, 661 (1991).

[5] R. Renner, Nature Phys. **3**, 645 (2007).

[6] R. Renner, Ph.D. thesis, ETH Zurich, 2005, arXiv:quant-ph/0512258;

[7] The permutation π on n elements acts on $\mathcal{H}^n = \mathcal{H}^{\otimes n}$ by permuting the tensor factors, i.e., $\pi|i_1\cdots i_n\rangle = |i_{\pi^{-1}(1)}\cdots i_{\pi^{-1}(n)}\rangle$ for a basis $\{|i\rangle\}$ of $\mathcal{H}$. The space of vectors invariant under the action of all π is denoted by $\mathrm{Sym}^n(\mathcal{H})$. As a map on $\mathrm{End}(\mathcal{H}^{\otimes n})$ we write $\pi(\rho) = \pi\rho\pi^{-1}$.

[8] The proof essentially relies on the fact that the state that maximizes the diamond norm can be post-selected by a measurement as constructed in Lemma 2.

[9] A purification $\tau_{\mathcal{H}^n\mathcal{R}}$ of $\tau_{\mathcal{H}^n}$ is a pure state on $\mathcal{H}^n\otimes\mathcal{R}$ satisfying $\mathrm{tr}_\mathcal{R}\tau_{\mathcal{H}^n\mathcal{R}} = \tau_{\mathcal{H}^n}$.

[10] K. Życzkowski and H.-J. Sommers, J. Phys. A **34**, 7111 (2001).

[11] M. Christandl, R. König, G. Mitchison, and R. Renner, Commun. Math. Phys. **273**, 473 (2007).

[12] Authenticity means that the communication cannot be altered by an adversary. If only completely insecure channels are available, authenticity may be simulated using a short initial key shared between Alice and Bob.

[13] This describes an entanglement-based protocol. However, our results immediately extend to prepare-and-measure schemes, because their security analysis can generally be reduced to corresponding entanglement-based schemes [14].

[14] C. H. Bennett, G. Brassard, and N. D. Mermin, Phys. Rev. Lett. **68**, 557 (1992).

[15] Of course, any nontrivial protocol generates keys of positive length $\ell > 0$ for at least some inputs.

[16] It is easy to see that this symmetrized key distillation protocol $\mathcal{E}$ satisfies $\mathcal{K}_\pi\circ\mathcal{E}\circ\pi = \mathcal{E}$ for any permutation π, where $\mathcal{K}_\pi$ is the operation that acts on the output (S_A, S_B, C) by replacing the communicated permutation $\bar{\pi}$ (in C) by $\bar{\pi}\circ\pi$. Similarly, we have $\mathcal{K}_\pi\circ(\mathcal{S}\circ\mathcal{E})\circ\pi = \mathcal{S}\circ\mathcal{E}$, because $\mathcal{K}_\pi$ acts like the identity on the key pair (S_A, S_B).

[17] The exponential de Finetti theorem provides an approximation of permutation-invariant states by states that are almost de Finetti states (i.e., convex combinations of states that have product structure on all but a few subsystems). Use of this result in a security proof will therefore—when compared to the case of collective attacks—lead to an extra additive error due to the approximation. In addition, it will require shortening the key because the approximation has the desired form only on almost all subsystems. The extra error and the amount by which the key has to be shortened obey a tradeoff. If, for instance, the additional error vanishes exponentially fast in n, then the key length must be decreased by an amount linear in n. In contrast, the work presented in this Letter does not result in an extra additive error but only a small (polynomial in n) multiplicative change in the security parameter [Eq. (4)], which can be compensated by shortening the key by an amount that is logarithmic in n.

[18] The role of $\mathcal{H}^{\otimes n}$ is taken by the space $\mathcal{V}$ of a finite-dimensional unitary representation V of G. We denote by $g|v\rangle \equiv V(g)|v\rangle$ the action of $g\in G$ on $|v\rangle\in\mathcal{V}$ and by $g(\rho) = V(g)\rho V(g^{-1})$ the action on $\mathrm{End}(\mathcal{V})$. The space $\mathcal{K}^{\otimes n}$ is replaced by a space $\mathcal{W}\cong\mathcal{V}$ on which G acts with the dual representation, $g|w\rangle = V(g^{-1})^T|w\rangle$ for $|w\rangle\in\mathcal{W}$. The role of the symmetric subspace $\mathrm{Sym}^n(\mathcal{H}\otimes\mathcal{K})$ is then taken by $(\mathcal{V}\otimes\mathcal{W})^G = \{|x\rangle\in\mathcal{V}\otimes\mathcal{W} : g\times g|x\rangle = |x\rangle\ \forall\ g\in G\}$, the invariant space of $\mathcal{V}\otimes\mathcal{W}$. The constant $g_{n,d}$ becomes $\dim(\mathcal{V}\otimes\mathcal{W})^G$ and the state $\tau_{\mathcal{V}\mathcal{W}}$ is the state proportional to the identity on $(\mathcal{V}\otimes\mathcal{W})^G\subset\mathcal{V}\otimes\mathcal{W}$.