Fully Device-Independent Quantum Key Distribution

Umesh Vazirani[1] and Thomas Vidick[2]
[1]*University of California, Berkeley, California 94720, USA*
[2]*California Institute of Technology, Pasadena, California 91125, USA*
(Received 19 June 2014; published 29 September 2014)

Quantum cryptography promises levels of security that are impossible to replicate in a classical world. Can this security be guaranteed even when the quantum devices on which the protocol relies are untrusted? This central question dates back to the early 1990s when the challenge of achieving device-independent quantum key distribution was first formulated. We answer this challenge by rigorously proving the device-independent security of a slight variant of Ekert's original entanglement-based protocol against the most general (coherent) attacks. The resulting protocol is robust: While assuming only that the devices can be modeled by the laws of quantum mechanics and are spatially isolated from each other and from any adversary's laboratory, it achieves a linear key rate and tolerates a constant noise rate in the devices. In particular, the devices may have quantum memory and share arbitrary quantum correlations with the eavesdropper. The proof of security is based on a new quantitative understanding of the monogamous nature of quantum correlations in the context of a multiparty protocol.

DOI: 10.1103/PhysRevLett.113.140501 PACS numbers: 03.67.Dd, 03.65.Ud, 03.67.Hk

The quest for unconditional security, or security based solely on the validity of quantum mechanics, is the holy grail of quantum cryptography. The rigorous proof that the BB84 [1] protocol for quantum key distribution (QKD) is unconditionally secure [2,3] (i.e., based solely on the validity of quantum mechanics) appeared to mark the successful end of this quest. Protocols for QKD allow two distant users, Alice and Bob, with the help of quantum devices, D_A and D_B, to establish a secret shared random key. Once a secret key has been generated the parties can exchange messages in a secure manner by using, e.g., a one-time pad. However, with the implementation of practical QKD systems came the realization that the quantum devices allowed for new kinds of attacks [4,5], for instance, where classical information gets leaked via quantum side channels [6]. Ruling out such attacks demands both a high level of technical ingenuity [7] as well as a measure of trust on the part of the user, thus calling into question the label of unconditional security.

Is there a principled way to rule out all such attacks? Mayers and Yao [8] were the first to put forth a challenge now known as device independence: Except for a necessary assumption of spatial separation, the quantum devices used would be treated as completely uncharacterized entities, and security would be guaranteed based solely on simple tests performed on the devices. That such a scheme for restoring unconditional security would even be possible relies on a unique feature of quantum entanglement, called monogamy [9]. Indeed, a hint of this approach can already be seen in Ekert's entanglement-based proposal for key distribution [10], which advocated tests based on the violation of Bell inequalities.

The first step toward a strong security guarantee was taken by Barrett *et al.* [11], who showed how to analyze a single round of interactions with the quantum devices. A sequence of follow-up papers extended these techniques to protocols where successive rounds are restricted to be completely independent [12–19] (i.e., the devices are memoryless), or causally independent [20,21]. In a recent breakthrough, two groups succeeded in designing secure protocols without making such independence assumptions [22,23]. However, the cost in obtaining such strong guarantees is that the analysis does not tolerate noisy devices, and the key rate, the number of bits of key extracted per use of the devices, tends to 0.

In this Letter we resolve the challenge of device-independent quantum key distribution (DIQKD) by showing that a variant of Ekert's original protocol has all the desirable features of DIQKD: It can be used to generate a shared random key at a 5% rate (30% of the raw key) while tolerating 2% noise rate in the devices (see Fig. 1 for a plot of the dependence of the key rate on the noise), and we establish its security against a general quantum eavesdropper. The security proof requires no independence assumptions—it only assumes that the devices can be modeled by the laws of quantum mechanics, and are spatially isolated from each other and from any adversary's laboratory. Our proof is thus the first to hold against the most general, "coherent" type of attacks. The key rate we obtain is nevertheless within a factor of 2 of the best rate known to be achievable even under much stronger assumptions, such as the assumption of individual attacks by the eavesdropper.

Since our security analysis ultimately rests on the violation of a Bell inequality, its successful experimental demonstration faces the same difficulties as every entanglement-based protocol for QKD. The most prominent hurdle is the closure of the so-called "detection loophole" [24], currently

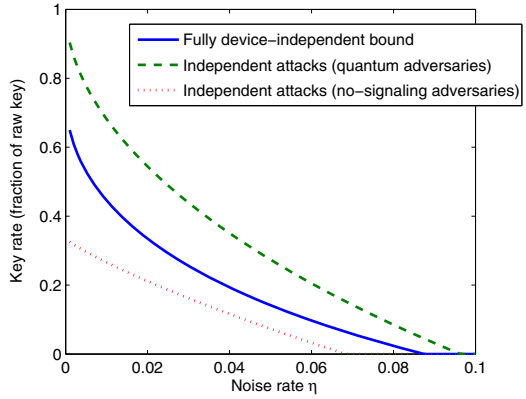

FIG. 1 (color online). Key rate obtained in our protocol (middle curve), expressed as a fraction of the raw key (bits obtained from the key rounds). On the x axis is the noise rate η as measured in the protocol. The top and bottom curves are the best achievable rates known for the case of quantum and no-signaling adversaries, respectively, under the additional assumption of causal independence [21].

a major challenge for experiments [25], but that has seen major advances in recent years [26,27]. The practical relevance of our results depends on, and we hope will help motivate, the resolution of this challenge. In our Letter we adopt a simple uniform model for the noise in which the measurement statistics obtained by honestly prepared devices are allowed to differ from the ideal statistics by a quantitative amount of at most $\eta \leq 2\%$, as measured by the statistical distance.

As previously mentioned, a property of quantum entanglement called monogamy plays a central role in the new DIQKD protocol and its analysis. In its simplest form the monogamy of entanglement states that two quantum systems that are maximally entangled cannot share any entanglement with a third system [28–32]. Intuitively, device-independent security can be derived from monogamy as follows: First, use the violation of a Bell inequality (such as the Clauser-Horne-Shimony-Holt [CHSH] inequality [33]) to establish that the correlations in the output distribution of the devices D_A and D_B are consistent with sharing a number of Einstein-Podolsky-Rosen (EPR) pairs $|\psi_{\mathrm{EPR}}\rangle = 1/\sqrt{2}(|00\rangle + |11\rangle)$. Next, conclude, based on monogamy of entanglement, that these correlations must be independent of any information that the eavesdropper, Eve, can obtain, thereby ensuring the security of the bits output by D_A and D_B. Crucially, this independence must be established even though D_A and D_B may have additional degrees of freedom that might be entangled with Eve's system.

The key insight of Barrett *et al.* [11,34] toward obtaining security guarantees for DIQKD was to focus on a weaker set of constraints than those imposed by quantum mechanics, specifically the no-signaling property dictated by special relativity. This line of attack provided a way of effectively quantifying, albeit only in the limited setting of a single round of interaction, the effects of monogamy at the level of the classical correlations between the systems. Extending this approach to realistic protocols with many rounds proved challenging [12–15], and lack of progress in this direction was recently explained in a beautiful result which establishes that approaches based solely on no-signaling cannot ultimately achieve security even for privacy amplification [35], a simpler task than QKD.

In contrast to the previous no-signaling approach [11,34], our analysis relies on a more complete picture of quantum mechanics. The key ingredient of quantum mechanics that we use is the existence of postmeasurement states. The properties of postmeasurement states play a crucial role in our analysis, in particular through the use of the so-called pretty-good measurement [36,37], a close-to-optimal distinguishing measurement defined directly from the postmeasurement states that plays a prominent role in quantum information theory. No analogue of this postmeasurement state exists for the case of purely nonsignaling theories.

The protocol and results.—Our DIQKD protocol is summarized in Table I. In each round $i \in \{1, \ldots, n\}$ of the protocol Alice selects a random trit x_i (element of $\{0, 1, 2\}$) as input to her device D_A and Bob selects a random bit y_i as input to his device D_B. Upon receiving their successive inputs, honest devices measure a fresh EPR pair in a particular choice of bases, returning outcomes a_i, b_i. Whenever $x \in \{0, 1\}$, the measurement bases are chosen to maximally violate the CHSH inequality,

$$\Pr(a_i \oplus b_i = x_i \wedge y_i) = \cos^2(\pi/8). \tag{1}$$

On the additional input $x = 2$, Alice's device D_A measures in the same basis as Bob's on input 1. This guarantees that $a_i = b_i$ whenever $x_i = 2$ and $y_i = 1$.

In the device-independent scenario, the devices are completely untrusted: For the analysis, the joint state of A, B and the eavesdropper E is modeled as an arbitrary quantum state ρ_{ABE} to which devices D_A and D_B apply arbitrary measurements at each stage of the protocol. The only assumption is that the three systems are localized: Device D_A (respectively, D_B) only has access to system A (respectively, B), and the eavesdropper to system E. In particular, the users' view of the protocol consists solely of the sequence of trits and bits they choose as inputs, and the sequence of bits they obtain in response. As we show, the phase of testing in step 3 of the protocol is sufficient to ensure security without further assumption about the inner workings of the devices.

Theorem.—Let p_s be the probability that the n-round protocol does not abort, when executed with devices D_A, D_B, a choice of noise tolerance $\eta = 0.02$ and a sufficiently small $\gamma > 0$. Let $\mathcal{E}$ be an arbitrary quantum system held by an eavesdropper, who in addition has access

TABLE I. Our DIQKD protocol requires the users, Alice and Bob, to make n uses of their devices. From the n pairs of output bits collected they are able to extract a shared key of length κn, where κ is a constant depending on the noise rate η that the users wish to tolerate.

1. Inputs: n = number of rounds, η = noise tolerance.
2. For rounds $i = 1, \ldots, n$: Alice picks $x_i \in \{0, 1, 2\}$, and Bob picks $y_i \in \{0, 1\}$, uniformly at random. They input x_i, y_i into their respective devices, obtaining outputs $a_i, b_i \in \{0, 1\}$ respectively.
3. Testing: Alice chooses a random subset $\mathbf{B} \subseteq \{1, \ldots, n\}$ of size γn, where γ is a small constant, and shares it publicly with Bob (rounds in $\mathbf{B}$ are called test rounds). Alice and Bob announce their input/output pairs in $\mathbf{B}$. They compute the fraction of inputs in $\mathbf{B}$ that satisfy the CHSH condition $a_i \oplus b_i = x_i \wedge y_i$. If this fraction is smaller than $\cos^2 \pi/8 - \eta$ they abort the protocol.
4. Extraction: Alice and Bob publicly reveal their choices of inputs. Let $\mathbf{C}$ be the set of rounds i in which $(x_i, y_i) = (2, 1)$ (rounds in $\mathbf{C}$ are called key rounds). The users compute the fraction of rounds in $\mathbf{B} \cap \mathbf{C}$ for which $a_i = b_i$. If it is less than $1 - \eta$ they abort the protocol. Otherwise, they perform information reconciliation on the remaining rounds in $\mathbf{C}$, followed by privacy amplification using, e.g., two-universal hashing.

to the classical information $\mathcal{P}$ exchanged by Alice and Bob on the authenticated public channel. Let A be the random variable describing the output of D_A, conditioned on the protocol not aborting. Then

$$\mathrm{H}^{\varepsilon}_{\min}(A|\mathcal{P}\mathcal{E}) \geq 0.05n - O(\ln(n/p_s\varepsilon)).$$

Here $\mathrm{H}^{\varepsilon}_{\min}$, the smooth conditional min entropy, is a measure of the correlations between the output of Alice's device and the eavesdropper's classical and quantum side information. Smooth conditional min entropy has been shown to be the appropriate measure of secrecy for the establishment of a universally composable secret key [38]. The theorem shows that even a fairly weak test (the CHSH game) between D_A and D_B strongly limits correlations between the eavesdropper's quantum state and the output of D_A. In the setting of our protocol, this establishes a very robust form of monogamy: It holds even in the presence of a substantial amount of leaked classical information $\mathcal{P}$ and assuming potentially extremely weak (p_s, ε can both be exponentially small in n) correlations between D_A and D_B.

More precisely, we obtain the following trade-off between the min entropy and the noise rate. Suppose that the protocol is run with devices such that the observed CHSH correlations in the test rounds satisfy $S = 2\sqrt{2}(1 - 2Q)$, where $2\sqrt{2}$ is the maximal possible violation of the CHSH inequality by quantum mechanics. Here Q is related to our "noise parameter" η by $Q = \eta/\sqrt{2}$. Then

$$\mathrm{H}^{\varepsilon}_{\min}(A|\mathcal{P}\mathcal{E}) \geq -\frac{11}{3}\log\left(\frac{11}{12} + \frac{2\sqrt{2}}{3}Q\right),$$

and after information reconciliation and privacy amplification the key rate r (expressed as a fraction of the key rounds $\mathbf{C}$) satisfies $r \geq H^{\varepsilon}_{\min}(A|\mathcal{P}\mathcal{E}) - h(Q)$, where h is the binary entropy function. Although somewhat worse than results obtained under the assumption of individual [13,14] or collective [15,21] attacks (see Fig. 1 for a comparison), our trade-off is already highly nontrivial and it is likely that further work will lead to improvements.

Security proof.—The goal of the security proof is to establish that the extracted key is random and secret under the sole assumption that the devices pass the testing phase of the protocol. Let n denote the number of rounds of the protocol, and let $A \in \{0, 1\}^n$ be the string of bits produced by Alice's device. The privacy amplification step at the end of the protocol ensures that to extract a secure key it is sufficient to bound Eve's information about A as $\mathrm{H}^{\varepsilon}_{\min}(A|E) \leq \alpha n$ for some $0 < \alpha < 1$. So we assume for contradiction that even though the devices pass the test described in the protocol (Table I) with non-negligible probability, Eve gains significant information about A. The goal is to use the success in the CHSH games between devices D_A and D_B to derive a contradiction. The main challenge is that the correlations between Eve and A are in fact very weak: Our only starting point is that she may gain non-negligible, but still potentially very small, information about a (possibly, again, quite small) part of the key. In addition, the tripartite interaction created by the protocol is complex and involves many rounds. This all but rules out the use of techniques, such as semidefinite programming [21] or an explicit modeling using linear algebra [19], that have been useful in simpler contexts.

We proceed in two steps. The first step relies on a powerful technique called the quantum reconstruction paradigm [39,40], which figured prominently in recent work on certifiable quantum randomness [41], a task originally proposed in [42] and studied further in [43]. Suppose that $\mathrm{H}^{\varepsilon}_{\min}(A|E) \leq \alpha n$. Then the reconstruction paradigm says that there exists a bit string Z (depending on A) of length roughly αn, such that given Z, Eve can choose a measurement of her quantum system E that results in a guess for A. Moreover, this guess will be correct with probability that scales polynomially with ε/n (see Lemma 3 in the Supplemental Material [44] for a precise statement). Crucially, in contrast to the direct interpretation of the (nonsmoothed) min entropy as a guessing probability [46], here the success probability does not depend on the initial uncertainty αn. The fact that Eve can now predict A, albeit based on additional "advice bits" in the form of the string Z, is essential for the remainder of our argument.

As a result we have obtained a stronger form of the adversary who is able to correctly guess, with small

probability, the whole output string A. Using Bayes' rule one can then derive the following:

$$\begin{aligned}&\Pr(\text{Eve guesses} A_1 \ldots A_n \text{correctly}) \geq \varepsilon \Rightarrow \exists i_0,\\ &\Pr(\text{Eve guesses} A_i | \text{guesses for}\\ &\quad A_1, \ldots, A_{i_0-1} \text{were correct}) \geq 1-\delta \end{aligned} \tag{2}$$

for some small $\delta > 0$. This concludes the first step of our proof: We used the quantum reconstruction paradigm to amplify the small, diffuse information that Eve may have had about the key bits into the existence of a measurement on her quantum system that successfully predicts, with high probability, the bit output by D_A in a single round i_0 of the protocol.

In the second step of the proof we derive a contradiction between the above strong form of the adversary and the CHSH tests performed as part of the protocol. To this end we introduce a simple argument, called the "guessing game," that is flexible and lends itself to many different types of scenario. In its basic form the guessing game states that, if two parties Alice and Bob are spatially separated and restricted to interacting with their own quantum devices, then Bob cannot obtain any information at all about Alice's input to her device. This is exactly a reformulation of the no-signaling condition. More precisely, we show that the first two of the following three conditions are incompatible with the third by showing that if both (a) and (b) held simultaneously in any single round, then the corresponding devices could be used to provide a successful (hence necessarily signaling) strategy in the guessing game: (a) the devices violate the CHSH inequality, whenever the round was selected as a test round; (b) the adversary can predict the output of Alice's device, whenever the round was selected as a key round; and (c) the no-signaling condition is satisfied between all three parties (Alice, Bob and the adversary).

Figure 2 gives a pictorial representation of a triple of devices that would satisfy all three conditions in a single round. We give a quantitative argument showing the precise tradeoff between them in Lemma 5 in the Supplemental Material [44]. Recall that as a consequence of the first step, we had identified a round i_0, at the start of which the two devices can be initialized in a state ρ_{i_0} [obtained through the conditioning on outcomes in previous steps that follows from our application of Bayes' rule (2)] such that all three conditions (a), (b), and (c) from the guessing game are satisfied. Indeed, as a result of the conditioning performed there exists a purification of ρ_{i_0} and a measurement for Eve on that purification that will produce the same outcome as Alice's device D_A, whenever its input is a 2. Moreover, when given inputs in $\{0,1\}$ the devices will produce outputs satisfying (1). Applying the guessing game leads to a contradiction with the no-signaling condition.

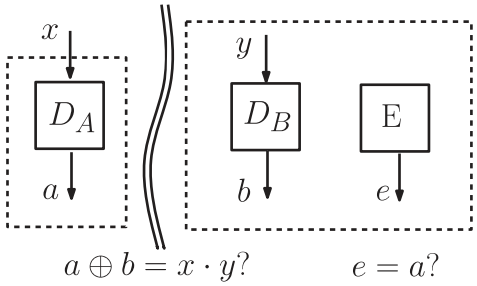

FIG. 2. The guessing game. Any devices satisfying both the CHSH condition $a \oplus b = x \cdot y$ and the guessing condition $a = e$ with high enough probability must allow signaling between D_A and $D_B + E$.

To conclude the proof we need to overcome a substantial difficulty that arises from the conditioning performed as part of the first step of the argument, in (2). Since Eve's measurement, as obtained from the quantum reconstruction paradigm, may depend on all public information in the protocol, including the users' choice of inputs as well as the information leaked through the advice bit string Z, by conditioning on the outcome of that measurement we are making a postselection on the shared quantum state of the devices D_A and D_B at the beginning of round i_0. The result is that condition (c) required for the application of the guessing game, while it automatically holds *a priori* (since the devices are nonsignaling by assumption), may no longer be satisfied once the conditioning has been performed. The difficulty is similar to one encountered in the analysis of the parallel repetition problem in two-player games [47], where conditioning on success in a subset of the repeated games may introduce correlations among the players in the remaining games. Here, the situation is further complicated by the presence of entanglement between the three parties and the fact that they are engaged in a relatively complex interaction.

To overcome this last difficulty we again formulate the situation as a game played between the different parties. The key idea is that, if it was the case that Alice and Bob's state depended on the choice of inputs made by Eve to perform her guessing measurement, then they could perform a measurement on that postmeasurement state to recover information about Eve's choice of inputs, thereby violating the basic no-signaling assumption. Formally, we obtain the following relation (Claim 9 in the Supplemental Material [44]):

$$\left(2n - \sum_i I(AB\!:\!X_i Y_i)_{\rho_i}\right) + \log(1/\varepsilon) \leq 2n, \tag{3}$$

where $I(AB\!:\!X_i Y_i)_{\rho_i}$ is the quantum mutual information between the quantum state of the two devices at the start of round i and inputs that were chosen for the devices in that round. The proof of Eq. (3) uses ideas originating in the coding strategy used in the proof of the

Holevo-Schumacher-Westmoreland theorem [48,49]. Combined with an application of the quantum Pinsker's inequality, Eq. (3) implies that, in an average round i, the state ρ_i is approximately independent from the inputs in round i, as required.

Discussion.—The protocol presented in this paper is the first major example of a purely classical tester for untrusted quantum devices that is provably robust against noise. Can our techniques be used more generally in the design of robust classical testers for untrusted quantum devices? One natural context to consider is certifiable quantum randomness [41], where the quantum reconstruction paradigm figured prominently in proving security against quantum adversaries. In an earlier version of our Letter we asked whether it is possible to design a protocol for randomness expansion that retains security against quantum adversaries while simultaneously being robust against noise. This question was recently answered affirmatively in [50].

More generally, with the maturing of quantum technology, there is intense interest in the question of how to test whether a quantum device behaves according to specification. Besides quantum cryptography, this issue arises in testing that a quantum computer is really quantum [51,52], and more generally in controlling the time evolution of an adversarial quantum system [23], and even testing the limits of quantum mechanics [53]. Solutions to date have to assume either that the tester can itself use some quantum resources, or that the protocol is entirely noise free: The only known proof [23] that an untrusted quantum system must evolve according to specification relies on directly characterizing the quantum state of any devices passing the tests of the classical tester, at the cost of a protocol that is not robust against noise. We are hopeful that our work will help open the way for the use of monogamy-based arguments toward the design of robust classical testers for untrusted quantum devices.

U. V. is supported by ARO Grant No. W911NF-12-1-0541, NSF Grant No. CCF-0905626, and Templeton Foundation Grant No. 21674. Part of this work was completed while T. V. was visiting UC Berkeley. T. V. is supported by the National Science Foundation under Grant No. 0844626 and by the Ministry of Education, Singapore under the Tier 3 Grant No. MOE2012-T3-1-009.

[1] C. Bennett and G. Brassard, in *Proceedings of IEEE International Conference on Computers, Systems, and Signal Processing, New York, 1984* (IEEE, Washington, DC, 1984), pp. 175–179.
[2] D. Mayers, J. ACM **48**, 351 (2001).
[3] P. W. Shor and J. Preskill, Phys. Rev. Lett. **85**, 441 (2000).
[4] C. H. Bennett, F. Bessette, G. Brassard, L. Salvail, and J. Smolin, J. Cryptol. **5**, 3 (1992).
[5] I. Gerhardt, Q. Liu, A. Lamas-Linares, J. Skaar, C. Kurtsiefer, and V. Makarov, Nat. Commun. **2**, 349 (2011).
[6] G. Brassard, N. Lütkenhaus, T. Mor, and B. C. Sanders, Phys. Rev. Lett. **85**, 1330 (2000).
[7] H.-K. Lo, X. Ma, and K. Chen, Phys. Rev. Lett. **94**, 230504 (2005).
[8] D. Mayers and A. Yao, in *Proceedings of the 39th Annual Symposium on Foundations of Computer Science, Palo Alto, 1998* (IEEE, Washington, DC, 1998), p. 503.
[9] B. Terhal, IBM J. Res. Dev. **48**, 71 (2004).
[10] A. K. Ekert, Phys. Rev. Lett. **67**, 661 (1991).
[11] J. Barrett, L. Hardy, and A. Kent, Phys. Rev. Lett. **95**, 010503 (2005).
[12] A. Acín, N. Gisin, and L. Masanes, Phys. Rev. Lett. **97**, 120405 (2006).
[13] A. Acín, S. Massar, and S. Pironio, New J. Phys. **8**, 126 (2006).
[14] V. Scarani, N. Gisin, N. Brunner, L. Masanes, S. Pino, and A. Acín, Phys. Rev. A **74**, 042339 (2006).
[15] L. Masanes, R. Renner, M. Christandl, A. Winter, and J. Barrett, arXiv:quant-ph/0606049v4.
[16] L. Masanes, Phys. Rev. Lett. **102**, 140501 (2009).
[17] E. Hänggi, R. Renner, and S. Wolf, in *Proceedings of the 29th EUROCRYPT, French Riviera* (Springer-Verlag, Berlin, 2010), pp. 216–234.
[18] A. Acín, N. Brunner, N. Gisin, S. Massar, S. Pironio, and V. Scarani, Phys. Rev. Lett. **98**, 230501 (2007).
[19] S. Pironio, A. Acín, N. Brunner, N. Gisin, S. Massar, and V. Scarani, New J. Phys. **11**, 045021 (2009).
[20] E. Hänggi and R. Renner, arXiv:1009.1833.
[21] L. Masanes, S. Pironio, and A. Acín, Nat. Commun. **2**, 238 (2011).
[22] J. Barrett, R. Colbeck, and A. Kent, Phys. Rev. A **86**, 062326 (2012).
[23] B. Reichardt, F. Unger, and U. Vazirani, Nature (London) **496**, 456 (2013).
[24] I. Gerhardt, Q. Liu, A. Lamas-Linares, J. Skaar, V. Scarani, V. Makarov, and C. Kurtsiefer, Phys. Rev. Lett. **107**, 170404 (2011).
[25] N. Gisin, S. Pironio, and N. Sangouard, Phys. Rev. Lett. **105**, 070501 (2010).
[26] B. G. Christensen, K. T. McCusker, J. B. Altepeter, B. Calkins, T. Gerrits, A. E. Lita, A. Miller, L. K. Shalm, Y. Zhang, S. W. Nam, N. Brunner, C. C. W. Lim, N. Gisin, and P. G. Kwiat, Phys. Rev. Lett. **111**, 130406 (2013).
[27] M. Giustina, A. Mech, S. Ramelow, B. Wittmann, J. Kofler, J. Beyer, A. Lita, B. Calkins, T. Gerrits, S. W. Nam, R. Ursin, and A. Zeilinger, Nature (London) **497**, 227 (2013).
[28] M. Fannes, J. Lewis, and A. Verbeure, Lett. Math. Phys. **15**, 255 (1988).
[29] R. F. Werner, Lett. Math. Phys. **17**, 359 (1989).
[30] V. Coffman, J. Kundu, and W. K. Wootters, Phys. Rev. A **61**, 052306 (2000).
[31] A. Higuchi, A. Sudbery, and J. Szulc, Phys. Rev. Lett. **90**, 107902 (2003).
[32] S. Bravyi, Quantum Inf. Comput. **5**, 216 (2005).
[33] J. F. Clauser, M. A. Horne, A. Shimony, and R. A. Holt, Phys. Rev. Lett. **23**, 880 (1969).
[34] J. Barrett, N. Linden, S. Massar, S. Pironio, S. Popescu, and D. Roberts, Phys. Rev. A **71**, 022101 (2005).
[35] R. Arnon-Friedman and A. Ta-Shma, Phys. Rev. A **86**, 062333 (2012).
[36] P. Hausladen and W. K. Wootters, J. Mod. Opt. **41**, 2385 (1994).

[37] H. Barnum and E. Knill, J. Math. Phys. (N.Y.) **43**, 2097 (2002).
[38] R. Renner, Ph.D. thesis, Swiss Federal Institute of Technology Zurich, 2005, arXiv:quant-ph/0512258.
[39] A. De and T. Vidick in *Proceedings of the 42nd ACM Symposium on Theory of Computing, Cambridge, MA, 2010* (ACM, New York, 2010), pp. 161–170.
[40] A. De, C. Portmann, T. Vidick, and R. Renner, SIAM J. Comput. **41**, 915 (2012).
[41] U. Vazirani and T. Vidick, in *Proceedings of the 44th Symposium on Theory of Computing, New York, 2012* (ACM, New York, 2011), pp. 61–76.
[42] R. Colbeck, Ph.D. thesis, Trinity College, University of Cambridge, 2006, arXiv:0911.3814.
[43] S. Pironio, A. Acin, S. Massar, A. B. De La Giroday, D. N. Matsukevich, P. Maunz, S. Olmschenk, D. Hayes, L. Luo, T. A. Manning *et al.*, Nature (London) **464**, 1021 (2010).
[44] See Supplemental Material at http://link.aps.org/supplemental/10.1103/PhysRevLett.113.140501, which includes Ref. [45].
[45] M. Wilde, *From Classical to Quantum Shannon Theory* (Cambridge University Press, New York, 2013).
[46] R. König, R. Renner, and C. Schaffner, IEEE Trans. Inf. Theory **55**, 4337 (2009).
[47] R. Raz, SIAM J. Comput. **27**, 763 (1998).
[48] A. Holevo, IEEE Trans. Inf. Theory **44**, 269 (1998).
[49] B. Schumacher, Phys. Rev. A **54**, 2614 (1996).
[50] C. A. Miller and Y. Shi, in *Proceedings of the 46th Annual ACM Symposium on Theory of Computing, Atlanta, 2014* (ACM, New York, 2014).
[51] A. Broadbent, J. Fitzsimons, and E. Kashefi, in *Proceedings of the 50th Annual Symposium on Foundations of Computer Science, New York, 2009* (IEEE, New York, 2009), pp. 517–526.
[52] D. Aharonov, M. Ben-Or, and E. Eban, in *Proceedings of the ICS, Beijing, China, 2010* (Tsinghua University Press, Beijing, 2010), pp. 453–469.
[53] D. Aharonov and U. Vazirani, in *Computability: Turing, Gödel, Church, and Beyond*, edited by B. Copeland, C. Posy, and O. Shagrir (MIT Press, Cambridge, MA, 2013).