PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A

Phil. Trans. R. Soc. A (2012) **370**, 3432–3448
doi:10.1098/rsta.2011.0336

Certifiable quantum dice

By Umesh Vazirani[1,*] and Thomas Vidick[2]

[1] *Computer Science Division, University of California at Berkeley, Berkeley, CA 94720-1776, USA*
[2] *Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA 02139, USA*

We introduce a protocol through which a pair of quantum mechanical devices may be used to generate n random bits that are ε-close in statistical distance from n uniformly distributed bits, starting from a seed of $O(\log n \log 1/\varepsilon)$ uniform bits. The bits generated are certifiably random, based only on a simple statistical test that can be performed by the user, and on the assumption that the devices obey the no-signalling principle. No other assumptions are placed on the devices' inner workings: it is not necessary to even assume the validity of quantum mechanics.

Keywords: randomness; Turing; entanglement; certifiable; extractors

1. Introduction

A source of independent and uniform random bits is a basic resource in many computational tasks, such as cryptography, game-theoretical protocols, algorithms and physical simulations. The quest for good hardware number generators goes as far back as the first commercially available computer, the Ferranti Mark I, a project in which Turing himself was very involved. That project was ultimately unsuccessful [1], and indeed constructing a physical source of randomness is an unexpectedly tricky task[1]—one that touches on fundamental questions about the nature of randomness. The focus of this paper is the fundamental philosophical issue: how can one certify whether one has succeeded? In other words, suppose someone were to claim that a particular device outputs uniformly random bits; is there a feasible test to verify that claim? In this paper, we report progress on this issue, in the form of a specific construction of a device for generating randomness, and an associated test that certifies (convincingly to a user who is afraid that the device was tampered with by an adversary) that the *particular* output produced by the device was drawn from a distribution that is statistically close to uniform.

*Author for correspondence (vazirani@cs.berkeley.edu).

[1]There has been extensive research into the design of practical hardware number generators. Intel recently announced the first *digital* such generator, as part of its new 'Ivy bridge' microprocessor [2].

One contribution of 18 to a Theme Issue 'The foundations of computation, physics and mentality: the Turing legacy'.

This might sound far fetched, because a uniform random number generator must output every n-bit sequence with equal probability $1/2^n$, and there seems to be no basis on which to reject any particular output in favour of any other. On the face of it, testing the output of the box amounts to classifying a *single* n-bit string x (or a very small sample of the exponentially many n-bit strings) as being random or not random. This naturally brings to mind the work of Kolmogorov, Chaitin and Martin-Löf on algorithmic randomness, which provides a foundation for probability theory based on a precise definition for what it means for a particular string x to be random: it is random if it is incompressible, or, as was shown to be essentially equivalent, it passes all computable statistical tests for randomness (we refer to the excellent book [3] for a comprehensive treatment of this aspect). Unfortunately, under this definition, whether or not x is random is in fact uncomputable, and even restricting the definition to efficient statistical tests (i.e. polynomial time) renders the randomness certification problem co-NP-complete. The theory of complexity-based pseudo-randomness suggests a way around this impasse, by showing that the outputs of certain cryptographically secure pseudo-random generators pass all polynomial time statistical tests (see Goldreich [4] for a recent survey). However, this very strong guarantee on the distribution of strings output by the generator is based on unproven cryptographic assumptions. Moreover, the generator only stretches randomness by a polynomial factor: to output an n-bit pseudo-random sequence, the generator must be fed n^c uniformly random bits, for some constant $c > 0$.

Starting in the mid-1980s, computer scientists explored a different approach to the question of designing a uniform random number generator. To avoid grappling with difficult philosophical questions about the nature of randomness, they assumed that they already had access to a physical device that was guaranteed to output random strings, except that the randomness was of 'low quality'. They modelled such devices as adversarially controlled sources of randomness, starting with the semi-random source [5], and weak random sources [6]. This sequence of papers has culminated in sophisticated algorithms called randomness extractors that are guaranteed to output a sequence that is arbitrarily close to truly random bits from physical sources of low-quality randomness (see Shaltiel [7] for a survey). It was clear, in a classical world, that these results were the best one could hope for—because it was necessary to assume that randomness in some form was output by the device in the first place, the only progress could be in minimizing the assumptions placed on the quality of that randomness.

Unlike classical physics, where randomness is implicitly an assertion about our lack of knowledge or computational ability, quantum mechanics offers a source of intrinsic randomness, ensconced in the Born rule, one of the fundamental axioms of the theory. So, in principle, it is very simple to design a quantum device that outputs a sequence of independent unbiased bits. In fact, there have been several concrete proposals for implementing random generators based on quantum mechanics [8,9], and testing whether the outputs of the generators pass certain simple statistical tests [10,11], or even tests specifically designed towards detecting algorithmic randomness [12]. Of course, success in passing a handful of such tests in no way certifies that the output distribution of the generator is close to uniformly random—the best that one can hope to say is that failure to pass the tests certifies that the generator's output distribution is not uniform.

This brings us back to the main question addressed in this paper: is it possible to positively certify that the output of a randomness-generating device (based on quantum mechanics) is 'really random' even though the user does not trust the experimental skills of the manufacturer, the calibration of the device, the manufacturer's motivations (particularly in cryptographic settings) or even the correctness of quantum mechanics. We describe the construction of a specific kind of quantum random number generator for which the answer to these questions is affirmative. This construction builds upon a proposal of Colbeck [13] and follow-up work by Pironio *et al.* [14] that provides a link between randomness certification and quantum non-locality. Before we can describe the actual construction, we must introduce some basic ideas behind quantum non-locality.

(*a*) *The Clauser–Horne–Shimony–Holt game*

Non-locality is one of the most interesting features of quantum mechanics and was explored in the famous work by Einstein, Podolsky and Rosen [15], and later in the work of Bell [16,17]. We focus here on a concrete realization of an experiment inspired by Bell's work that is best phrased as a game, the Clauser–Horne–Shimony–Holt (CHSH) game, named after its inventors Clauser *et al.* [18]. The CHSH game is played by two cooperating players, Alice and Bob, who we model using two spatially separated boxes $\mathcal{A}$ and $\mathcal{B}$. Each of the boxes may contain parts of an entangled quantum system,[2] and, in general, they can be described through the underlying probability distribution $p_{\mathcal{AB}}(a, b|x, y)$, the probability of the boxes producing outputs a and b when provided inputs x and y, respectively.

The spatial separation between the boxes reflects the fact that Alice and Bob are not allowed to communicate during the game. Mathematically, this assumption is enforced through a *no-signalling* condition placed on the distribution $p_{\mathcal{AB}}$. This condition states that the marginal distribution on $\mathcal{A}$'s outputs (respectively, $\mathcal{B}$'s outputs) should be independent of $\mathcal{B}$'s input (respectively, $\mathcal{A}$'s input),

$$\forall a, x, y, y', \quad \sum_b p_{\mathcal{AB}}(a, b|x, y) = \sum_b p_{\mathcal{AB}}(a, b|x, y'), \tag{1.1}$$

and a symmetric equation should hold when summing over a.

At the start of the game, each player is given an input $x, y \in \{0, 1\}$ chosen uniformly at random. They are allowed to perform arbitrary measurements on their share of the entangled state, but are not allowed to communicate. Their goal is to produce outputs a, b that satisfy the *CHSH condition*

$$a \oplus b = x \wedge y. \tag{1.2}$$

It is not hard to see that, if the boxes are only classically correlated (i.e. their only non-local resource is shared randomness), the best strategy for the players will lead to a success probability of $\frac{3}{4}$—in fact, systematically outputting 0 is the best one can do. However, there are quantum measurements on a

[2]For the purposes of this introductory discussion, we assume that the players in the game obey the laws of quantum mechanics, although that assumption will be relaxed later on.

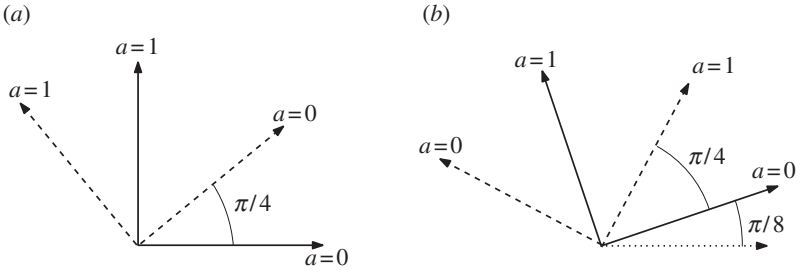

Figure 1. The bases used in the CHSH game. Solid lines correspond to the basis used on input $x=0$ (a) and $y=0$ (b). Dotted lines correspond to the basis used on input $x=1$ (a) and $y=1$ (b). Pairs of vectors corresponding to valid outputs always make an angle of $\pi/8$.

2-qubit entangled state that allow the players to obtain a strictly higher success probability of $\cos^2 \pi/8 \approx 0.85$. The entangled state is a *Bell pair*

$$|\Psi\rangle = \frac{1}{\sqrt{2}}(|0\rangle|0\rangle + |1\rangle|1\rangle).$$

This notation describes the joint state of a pair of two-dimensional systems that are in an equal superposition of two identical states: the $|0\rangle|0\rangle$ state and the $|1\rangle|1\rangle$ state. Each party will measure its own half of $|\Psi\rangle$ using one of the two possible choices of basis, depending on the input bit. These bases are such that, out of the four pairs of bases, those corresponding to input pairs (x, y) such that $x \wedge y = 0$ make an angle $\pi/8$ with each other, while the pair corresponding to the inputs $(1, 1)$ make an angle $\pi/2 - \pi/8$ (see figure 1 for an illustration). The resulting measurements have the property that, for *every* possible pair of inputs, the pair of outputs obtained by making the corresponding measurements on both halves of $|\Psi\rangle$ will be correct with probability $\cos^2 \pi/8$.

In fact, for any number $\frac{3}{4} < p_{\mathsf{CHSH}} \leq \cos^2 \pi/8 \approx 0.85$, there is a quantum strategy achieving exactly that success probability. Hence, we may define the *quantum regime* for the CHSH game as this range of probabilities: for any value in that range, there is a simple quantum-mechanical pair of boxes, whose underlying distribution $p_{\mathcal{AB}}$ obeys the no-signalling condition (1.1), which achieves that success probability.

These well-known facts have a striking consequence, first made explicit in Colbeck's PhD thesis [13] (see also Colbeck & Kent [19] for an expanded version): any boxes producing correlations that fall in the quantum regime *must generate randomness*![3] Indeed, deterministic boxes—boxes whose underlying distribution $p_{\mathcal{AB}}$ is concentrated on a unique pair of outputs (a, b) for every possible input pair (x, y)—are inherently classical, so that their success probability must fall in the classical regime $p_{\mathsf{CHSH}} \leq \frac{3}{4}$. This observation provides a first glimpse of a possible test for randomness. The test would be based on the use of two spatially separated boxes. Inputs to the boxes would be chosen independently and uniformly at random, and their outputs would be checked for the CHSH constraint (1.2). Repeating the experiment many times, and verifying that the CHSH constraint

[3]The idea that the sole violation of a Bell inequality may be useful to cryptographic tasks by itself is not new, and dates back at least to the first proposals for device-independent quantum key distribution [20,21].

is satisfied more than 75 per cent of the time, would guarantee that the sequence of outputs produced by the boxes contains (some) randomness. Moreover, the certified presence of randomness would not depend on any assumption on the physical nature of the boxes—it is guaranteed by a simple statistical test, together with the no-signalling condition (1.1).

(*b*) *A randomness-expansion protocol*

The main drawback of the procedure described above is that in order to test the device one needs more randomness than is actually produced. Indeed, selecting inputs to the boxes in the CHSH game requires two bits of randomness, while certainly no more than two bits will be produced. Hence, while the bits output are certifiably random, one needs to start with an equal, if not larger, number of random bits prior to the experiment. In a recent paper published in *Nature*, Pironio *et al.* [14] showed that one could reliably (up to some error ε) certify that the average value of p_{CHSH} across n steps of interaction with the boxes is higher than $\frac{3}{4}$ by choosing inputs to the boxes according to a very biased distribution. n pairs of inputs following this biased distribution may be sampled using only $\sqrt{n \log 1/\varepsilon}$ uniformly random bits, resulting in a protocol that uses approximately $\sqrt{n \log 1/\varepsilon}$ random bits to certifiably generate n bits that are ε-close in statistical distance to being uniformly distributed.[4] (An extractor could then be applied to produce linearly many near-uniform bits.) Pironio *et al.* [14] also reported an experimental realization of their scheme, demonstrating the generation of 42 new random numbers, in addition to the randomness used to execute the protocol.

Ideally, one would like a randomness-generating device that requires very few or no random bits as input to get started. Indeed, a small amount of randomness is necessary, because if the tests were totally deterministic they would be passed by a box that outputs a predetermined sequence of bits. Quantitatively, one can argue that $\log 1/\varepsilon$ bits are necessary to guarantee that the output is ε-close to uniform in statistical distance. This is because a malicious device could attempt to guess the experimenter's private random bits, and behave accordingly. If the experimenter uses only s random bits, the device's guess will be successful with probability 2^{-s}, and, in that case, it can deterministically satisfy all the experimenter's requirements (since they are known in advance). One can also argue (perhaps a little less emphatically) that $\log n$ bits are necessary, since at best the device acts like a weak random source, and a random seed of $\log n$ bits is necessary to extract that randomness.

In this paper, we describe a very simple protocol that comes close to achieving this, using a random seed of length $O(\log n \log 1/\varepsilon)$ to generate an n-bit string that is ε-close to uniformly random in statistical distance. This exponentially improves upon the quadratic expansion of Pironio *et al.* [14]. The protocol to achieve this prescribes the interaction of a trusted user with an *untrusted physical device*, which we assume is made of two separate boxes, $\mathcal{A}$ and $\mathcal{B}$. The protocol consists of $m = O(n)$ phases. Each phase lasts for $k = O(\log n + \log 1/\varepsilon)$ rounds, during each of which the user inputs a single bit to each box and collects a single bit as output. This sequence of mk interactions is non-adaptive: the user can generate the $2mk$ input bits in advance from his $O(\log n \log 1/\varepsilon)$-bit random seed

[4]The paper [14] contained an error that was later fixed in the work of Fehr *et al.* [22].

before the interaction. It is of critical importance that he reveals the input bits of phase $i+1$ only after the completion of phase i. In each phase, the user also performs a very simple statistical test (which simply checks that the fraction of rounds that satisfy the CHSH condition lie in the quantum range, as described in §1*a*). If the test is passed in all phases, then the output of box $\mathcal{B}$, say, is efficiently (and classically) post-processed to output an n-bit string. If the test is failed in any phase, then the user outputs a special 'fail' symbol. (We refer to §2 for a more detailed description of the protocol and the relevant parameters.)

The bits produced in the protocol are guaranteed to be statistically close to uniformly distributed provided the following conditions are met:

— the user's private random bits are uniformly random;
— the simple statistical test is passed in all m phases of the protocol; and
— there is no communication between $\mathcal{A}$ and $\mathcal{B}$ in the middle of any phase. Formally, this requirement states the following: for each phase i, the marginal distribution of outputs produced by $\mathcal{A}$ (respectively, $\mathcal{B}$) during phase i is independent of all inputs to $\mathcal{B}$ (respectively, $\mathcal{A}$) during phase i.

As we shall see, it is possible to implement using very simple quantum mechanics a pair of boxes $\mathcal{A}$ and $\mathcal{B}$ that pass the statistical tests in all m phases with high probability. Equally interesting is the fact that the three conditions listed above made no mention of quantum mechanics. The main condition is that there must be no communication between the two boxes during the execution of a single phase. There are many ways one could imagine enforcing this condition, most markedly by ensuring that they are separated by a distance greater than that light can travel during each time interval. In this sense, the randomness in the output could be said to be 'Einstein certified'—it is certifiable even to a quantum-sceptic (Einstein's oft repeated quote 'God does not play dice with the Universe' comes to mind), as it only relies on special relativity together with a simple statistical test on the actual output of the device.

The proof of the guarantee on the randomness of the output of the boxes outlined above is based on a simple *guessing game* introduced in §3. In the guessing game, Bob is given an input $y \in \{0,1\}$, and Alice is asked to produce a guess for that input. Clearly, any strategy with success probability larger than $\frac{1}{2}$ (over the choice of Bob's input) indicates communication between Alice and Bob. Our analysis proceeds by showing that a pair of boxes $\mathcal{A}$ and $\mathcal{B}$ that pass the simple statistical test in every phase while not outputting sufficient randomness can actually be used to devise a successful strategy in the guessing game, thus contradicting the assumption that the boxes did not communicate in the middle of any phase.

(*c*) *Classical and quantum adversaries*

Among the many applications for random bits, some of the most prominent pertain to the area of cryptography. For instance, the most widely studied key distribution protocol, BB84 [23], requires a large number of uniformly distributed bits in order to make an initial choice of basis. In such an application, it is crucial that the bits generated appear close to uniform not only to the (honest) user of the cryptographic protocol, but also to any external *adversary* to the protocol.

A limited type of adversary is a *classical* adversary, who may share classical correlations with the device used in order to generate the random bits. Such an adversary can, for instance, be used to model a situation in which the device has been prepared by a malicious opponent using an arbitrary procedure. The classical correlations correspond to the information that the adversary has kept about the workings of the device. Recently, Pironio & Massar [24] and independently Fehr *et al.* [22] extended the results of Pironio *et al.* [14] to show that the same protocol could be used to produce bits secure against such classical adversaries.

A stronger type of adversary is a *quantum* adversary. Such an adversary can potentially be much more powerful than a classical adversary. For instance, she could in addition have decided to initialize the device in an entangled quantum state that extends into her own laboratory. This opens the possibility that the adversary may be able to perform specific measurements on her share of the entanglement, producing outcomes that are correlated with whatever bits the device may have produced in the protocol.

A reason to doubt that the adversary may gain much information in this way comes from a delicate property of entanglement, its *monogamy* [25]. Informally, monogamy states that a tripartite entangled state $|\Psi\rangle_{ABE}$ cannot be maximally entangled both between A and B and between B and E. Since our protocol enforces very strict correlations between the outputs of $\mathcal{A}$ and $\mathcal{B}$, one may hope that these correlations will pre-empt any strong correlation between either of them and an arbitrary adversary.

In Vazirani & Vidick [26], we show that this scenario can indeed be ruled out by proving an analogue to theorem 4.1, which also holds in the presence of a quantum adversary. The theorem applies to a slight variant of the protocol described in figure 2, based on using an 'extended' version of the CHSH game with four possible inputs per box. It also requires an additional assumption on the boxes, in addition to the no-signalling condition: that their inner workings may be described by quantum mechanics. Such an assumption is necessary if we wish to quantify the correlations between the boxes and a quantum adversary; however, the absence of such a requirement is one of the strengths of the procedure described in this paper.

2. Exponential expansion of randomness

Let n and ε be two fixed parameters: n is the target number of random bits that the user would like to generate at the end of the protocol, and $\varepsilon > 0$ is an 'error' parameter bounding the statistical distance between the output bits eventually produced and the uniform distribution on n bits.

Our protocol, summarized in figure 2, uses two main ideas in order to save on the randomness required for the user to execute it. The first idea is to restrict the inputs to $(0,0)$ most of the time. Only a few randomly placed checks (the Bell rounds) are performed in order to verify that the boxes are generating their inputs honestly. There are about $O(\log 1/\varepsilon)$ such blocks. Note that the boxes usually do not know when they are being checked: for instance, if the input is $(0,1)$ then, even though box $\mathcal{B}$ knows that it is in a Bell round, box $\mathcal{A}$ by itself cannot differentiate that particular round from one in which both inputs

protocol A

1. Let n, ε be parameters given as input. Set $m = 400n$, $\Delta = 100\lceil \ln(1/\varepsilon) \rceil$, $\ell = m/\Delta$ and $k = 100\lceil \log n + \log 1/\varepsilon \rceil$.
2. Choose $T \subseteq [m]$ uniformly at random by selecting each position independently with probability $1/\ell$.
3. Repeat, for $i = 1, \ldots, m$:
   - 3.1 if $i \notin T$, then
     - 3.1.1 set $x = y = 0$ and choose x, y as inputs for k consecutive steps. Collect outputs $a, b \in \{0,1\}^k$.
     - 3.1.2 if $a \oplus b$ has more than $\lceil 0.2k \rceil$ 1's then reject and abort the protocol. Otherwise, continue.
   - 3.2 if $i \in T$,
     - 3.2.1 pick $x, y \in \{0,1\}$ uniformly at random, and set x, y as inputs for k consecutive steps. Collect outputs $a, b \in \{0,1\}^k$.
     - 3.2.2 if $a \oplus b$ differs from $x \wedge y$ in more than $\lceil 0.2k \rceil$ positions then reject and abort the protocol. Otherwise, continue.
4. If all steps accepted, then accept.

Figure 2. Protocol A uses $O(\log n \log 1/\varepsilon)$ bits of randomness and makes $O(n(\log n + \log 1/\varepsilon))$ uses of the boxes. Theorem 4.1 (in §4) shows that n bits of randomness are produced with confidence $1 - \varepsilon$. The threshold $0.2k$ in steps 3.1.2 and 3.2.2 is arbitrary, and any value strictly lower than $k/4$ would work.

are 0. This implies, in particular, that the strategy it uses to determine its output cannot be different from what it would have been had the inputs been the more frequent $(0,0)$.[5]

The second main idea consists of decomposing the protocol into *blocks*, which are consecutive sequences of a fixed number $k = O(\log^2 n)$ of rounds of interaction between the user and the boxes. Each box always receives identical inputs throughout all rounds of a given block. The blocks' purpose is to enable the user to perform a robust verification of the CHSH condition: his final test will enforce that, in *every* block, a significantly larger than $\frac{3}{4}$ fraction of pairs of outputs satisfy the CHSH condition (with respect to the corresponding pair of inputs). This lets us argue about the *Hamming distance* between Alice and Bob's k-bit outputs in any block: if the CHSH constraint was of the form $a \oplus b = 0$ then the outputs should be close in Hamming distance, whereas if it had the form $a \oplus b = 1$ then they should be far apart.

Altogether, protocol A requires only the use of random bits in order to select the position of the Bell blocks, as well as to select inputs in these blocks. The $O(\log 1/\varepsilon)$ Bell blocks can be chosen among the $O(n)$ rounds using $O(\log n \log 1/\varepsilon)$ random bits [27], and corresponding uniformly distributed inputs may be generated using an additional $O(\log 1/\varepsilon)$ random bits.

[5]This idea was already used in Pironio *et al.* [14], and led to their protocol with quadratic $\sqrt{n \log 1/\varepsilon} \to n$ expansion of randomness.

Before proceeding, we should ensure that 'honest' boxes, which play the optimal quantum strategy for the CHSH game (as described in figure 1) independently in every round, are accepted by the user with very high probability. Indeed, we have seen that such boxes will satisfy the CHSH constraint independently with probability $\cos^2 \pi/8$ in each round. Hence, when one considers a block of k successive rounds, the probability that the CHSH constraint is *not* satisfied in more than 20 per cent of those rounds will be exponentially small in k. Precisely, a simple Chernoff bound shows that the probability that the honest strategy satisfies the CHSH condition in less than 80 per cent of any k successive rounds is at most $\exp(-2(\cos^2 \pi/8 - 0.80)^2 k) = \mathrm{e}^{-\Omega(k)}$. Given our choice of $k = 100\lceil \log n + \log 1/\varepsilon \rceil$, it can be verified that, for large enough n, this expression is smaller than ε/m, where $m = 400n$ is the total number of blocks in the protocol. By a union bound, such boxes will fail to produce correlations satisfying the user in even just one of these blocks with probability at most ε.

The key part of the proof of our result consists of analysing the case of *arbitrary* non-signalling boxes: we need to show that any such boxes that are accepted in the protocol *must* produce n bits of randomness. This requires that we define more precisely what we mean by 'n bits of randomness'. Recall that our ultimate goal is to produce bits that are uniformly random, as this will make them suitable for most applications. We will achieve this in two stages. In the first stage, we will show that the bits produced by one of the boxes, say $\mathcal{B}$, in the protocol must contain a linear amount of *smooth min-entropy*. In the second stage, we will observe that this condition is enough to guarantee that $\mathcal{B}$'s outputs may then be massaged into bits that are close to being uniformly distributed.

The min-entropy measures the highest 'peak' of a distribution: a random variable with min-entropy k never takes any given value with probability larger than 2^{-k}. Given a random variable B, it is defined as $H_\infty(B) := -\log \max_b \Pr(B = b)$.[6] The smooth min-entropy is a 'robust' variant of the min-entropy, defined as

$$H_\infty^\varepsilon(B) = \sup_{X, \|X-B\|_1 \leq \varepsilon} H_\infty(X),$$

where $\|X - B\|_1 := (\frac{1}{2}) \sum_b |\Pr(B = b) - \Pr(X = b)|$ is the statistical distance and $\varepsilon > 0$ is a 'smoothness' parameter.

It turns out that the choice of the smooth min-entropy is the right one with respect to the task of producing bits that are indistinguishable from uniform. Indeed, one of the major achievements of the field of randomness extraction is the demonstration that a random variable B with smooth min-entropy $H_\infty^\varepsilon(B) \geq K$ may be efficiently transformed into a string of approximately K bits that are ε-close to being uniformly distributed. This task can be accomplished using a combinatorial construction known as an extractor. An extractor usually takes two inputs. The first, called the *source*, is the random variable B from which we wish to extract uniformly random bits. The second is a random variable Y called the *seed*. It can be much shorter than the source, but is required to be uniformly

[6]The min-entropy is a stronger notion of entropy than the Shannon entropy, in the sense that having large min-entropy implies having large Shannon entropy, but the converse is not true.

distributed. The extractor maps the pair (X, Y) to a single random variable Z, over K bits, with the promise that Z is statistically close to uniform as long as B has at least K bits of min-entropy. In order to do so, the best extractors require the seed to be of length about $O(\log n + \log 1/\varepsilon)$ [28].

To summarize, our protocol will use a source of fresh random bits in each of its two stages. In the first stage, $O(\log n \log 1/\varepsilon)$ bits will be used in order to select roughly $O(n \log^2 n)$ pairs of inputs to the boxes, as described in figure 2. The user will then collect as many pairs of outputs, and verify that they satisfy the CHSH correlations. If not, he will reject the outputs. If so, he will use an additional $O(\log n + \log 1/\varepsilon)$ uniform bits to play the role of the seed Y in a specific extractor construction. He will apply the extractor to the bits B produced by the box $\mathcal{B}$ and Y, producing a string Z of n bits. The final output of the protocol is Z itself. Our main result, theorem 4.1 below, guarantees that Z is statistically close to a uniformly distributed n-bit string.

3. The guessing game

The analysis of our protocol is based on the definition of a simple 'guessing game'. In this game, there are two cooperating players, Alice and Bob. At the start of the game, Bob receives a single bit $y \in \{0, 1\}$ chosen uniformly at random. The players are then allowed to perform arbitrary computations, but are not allowed to communicate. At the end of the game, Alice outputs a bit a, and the players win if $a = y$.

Clearly, any strategy with success probability larger than $\frac{1}{2}$ indicates a violation of the no-communication assumption between Alice and Bob. At the heart of the proof of theorem 4.1 is a reduction to the guessing game. Assuming that there existed a pair of boxes violating the conclusions of the theorem, we will show how these boxes may be used to devise a successful strategy in the guessing game, contradicting the no-signalling assumption placed on the boxes.

To illustrate the main features of the strategies we will design later, consider the following simplified setting. Let $\mathcal{A}, \mathcal{B}$ be a given pair of boxes taking inputs $X, Y \in \{0, 1\}$ and producing outputs $A, B \in \{0, 1\}^k$, respectively. Assume the following two properties hold. First, if the input to $\mathcal{B}$ is $Y = 0$, then its output B is essentially deterministic, in the sense that $B = b_0$ with high probability. Second, whatever their inputs, the boxes' outputs satisfy the CHSH constraint on average with a slightly higher probability than could any classical boxes: there is a fixed $\delta > 0$ such that a fraction at least $\frac{3}{4} + \delta$ of $i \in [k]$ are such that $A_i \oplus B_i = X \wedge Y$. Then we claim that there is a strategy for Alice and Bob in the guessing game, using $\mathcal{A}$ and $\mathcal{B}$, that succeeds with probability strictly larger than $\frac{1}{2}$, demonstrating that the boxes must be signalling.

Alice and Bob's strategy is the following. Alice is given access to $\mathcal{A}$ and Bob to $\mathcal{B}$. Upon receiving his secret bit y, Bob inputs it to $\mathcal{B}$, collecting outputs $b \in \{0, 1\}^k$. Alice chooses an $x \in \{0, 1\}$ uniformly at random, and inputs it to $\mathcal{A}$, collecting outputs $a \in \{0, 1\}^k$. Let b_0 be the k-bit string with the highest probability of being output by $\mathcal{B}$, conditioned on $y = 0$. Alice makes a decision as follows: she computes the relative Hamming distance $d = d_{\mathrm{H}}(a, b_0)$. If $d < \frac{1}{4}$ she claims 'Bob's input was 0'. Otherwise, she claims 'Bob's input was 1'.

By assumption, if Bob's secret bit was $y = 0$, then his output is almost certainly b_0. By the CHSH constraint, independently of her input Alice's output a lies in a Hamming ball of radius $\frac{1}{4} - \delta$ around b_0. So in this case she correctly decides to claim 'Bob's input was 0'.

In the case that Bob's secret bit was $y = 1$, the analysis is more interesting. Let b be the actual output of $\mathcal{B}$. Let a_0 and a_1 be $\mathcal{A}$'s output in the two cases $x = 0$ and $x = 1$, respectively. We claim that the Hamming distance $d_{\mathrm{H}}(a_0, a_1) \geq \frac{1}{2} + 2\delta$. This is because, by the CHSH constraint, $d_{\mathrm{H}}(a_0, b) \leq \frac{1}{4} - \delta$, while $d_{\mathrm{H}}(a_1, b) \geq \frac{3}{4} + \delta$. Applying the triangle inequality

$$d_{\mathrm{H}}(a_0, a_1) \geq |d_{\mathrm{H}}(a_1, b) - d_{\mathrm{H}}(a_0, b)| \geq \tfrac{1}{2} + 2\delta,$$

as claimed. Hence both a_0 and a_1 cannot lie in the Hamming ball of radius $\frac{1}{4}$ around the fixed string b_0 (observe that this argument makes no use of the actual location of b_0!). Thus, in the case $y = 1$, Alice correctly outputs 'Bob's input was 1' with probability at least $\frac{1}{2}$.

Overall, Alice and Bob succeed in the guessing game with probability $\frac{3}{4}$, implying the boxes $\mathcal{A}$, $\mathcal{B}$ allowed them to communicate, and hence do not satisfy the no-signalling condition.

Clearly, there is a lot of slack in the above reasoning, because for contradiction it suffices to succeed in the guessing game with any probability strictly greater than $\frac{1}{2}$. By being more careful, it is possible to allow Bob's output on $y = 0$ to have more min-entropy, as well as allow for a small probability that the boxes' outputs may not satisfy the CHSH constraint:

Lemma 3.1. *Let $\beta, \gamma > 0$ be such that $\gamma/2 + 2\beta < \frac{1}{4}$, and k an integer. Suppose given a pair of boxes $\mathcal{A}, \mathcal{B}$, taking inputs $X, Y \in \{0,1\}$ and producing outputs $A, B \in \{0,1\}^k$ each. Suppose the following conditions hold:*

1. *When given input 0, the distribution of outputs of $\mathcal{B}$ has low min-entropy: there exists a $b_0 \in \{0,1\}^k$ such that $\Pr(B = b_0 | Y = 0) \geq 1 - \gamma$.*
2. *The boxes' outputs fall in the 'quantum regime' of the CHSH inequality: there exists a constant $\delta > 0$ such that*

$$\Pr(d_H(A \oplus B, (X \wedge Y, \ldots, X \wedge Y)) > \tfrac{3}{4} + \delta) \leq \beta,$$

where the probability is taken over the choice of uniformly random X, Y, and the boxes' internal randomness.

Then there is a strategy for Alice and Bob, using $\mathcal{A}$ and $\mathcal{B}$, which gives them success probability strictly greater than $\frac{1}{2}$ in the guessing game.

Proof. Alice and Bob's strategy in the guessing game is as described above. Let b_0 be the k-bit string that is most likely to be output by $\mathcal{B}$, conditioned on $y = 0$.

We first show that, if Bob's input was $y = 0$, then Alice claims that Bob had a 0 with probability at least $1 - \gamma - 2\beta$. By the first condition in the lemma, Bob obtains the output b_0 with probability at least $1 - \gamma$. Moreover, by the second condition, the CHSH constraint will be satisfied with probability at least $1 - 2\beta$ on average over Alice's choice of input, given that Bob's input was $y = 0$. Given $y = 0$,

whatever the input to $\mathcal{A}$, the CHSH constraint implies that $d_{\mathrm{H}}(a, b) < \frac{1}{4}$. Hence by a union bound Alice will obtain an output string a at relative Hamming distance at most $\frac{1}{4}$ from b_0 with probability at least $1 - \gamma - 2\beta$.

Next we show that, in the case where Bob's input in the guessing game is $y = 1$, Alice claims that Bob had a 1 with probability at least $\frac{1}{2}(1 - 8\beta)$. The second condition in the lemma implies that, for any of the two possible choices for Alice's input $X = x \in \{0, 1\}$, it holds that

$$\Pr_{ABY}(d_{\mathrm{H}}(A \oplus B, (X \wedge Y, \ldots, X \wedge Y)) > \tfrac{3}{4} + \delta | X = x) \leq 2\beta. \tag{3.1}$$

Let b' be Bob's output, and suppose that b' is such that, for every $x \in \{0, 1\}$,

$$\Pr_{AY}(d_{\mathrm{H}}(A \oplus B, (X \wedge Y, \ldots, X \wedge Y)) > \tfrac{3}{4} + \delta | X = x, B = b') \leq 4\beta$$

holds. It follows from (3.1) and Markov's inequality that this condition holds with probability at least $1 - 4\beta$ over b'.

If Alice chooses $x = 0$, then the CHSH constraint indicates that the corresponding a_0 should be such that $d_{\mathrm{H}}(a_0, b') \leq \frac{1}{4} - \delta$, while in the case that she chooses $x = 1$ her output a_1 should satisfy $d_{\mathrm{H}}(a_1, b') \geq \frac{3}{4} + \delta$. By the triangle inequality,

$$d_{\mathrm{H}}(a_0, a_1) \geq |d_{\mathrm{H}}(a_1, b') - d_{\mathrm{H}}(a_0, b')| \geq \tfrac{1}{2} + 2\delta,$$

so that, whatever the value of b', at most one of a_0 or a_1 can be at distance less than $\frac{1}{4}$ from b_0. Because Alice's input is chosen uniformly at random, taking into account the choice of b' we have shown that with probability at least $(1 - 4\beta)/2$ Alice will choose an input that will make her correctly claim that Bob had a 1.

The two bounds proved earlier together show that Alice's probability of correctly guessing Bob's input in the guessing game is at least

$$p_{\mathrm{succ}} \geq \frac{1}{2}(1 - \gamma - 2\beta) + \frac{1}{2}\frac{1 - 4\beta}{2} = \frac{1}{2} + \left(\frac{1}{4} - 2\beta - \frac{\gamma}{2}\right),$$

which is greater than $\frac{1}{2}$ whenever $2\beta + \gamma/2 < \frac{1}{4}$, proving lemma 3.1. ■

4. Analysis

In this section, we state and prove our main theorem.

Theorem 4.1. *Let $\varepsilon > 0$ be given, and n an integer. Let $(\mathcal{A}, \mathcal{B})$ be an arbitrary pair of no-signalling boxes used to execute protocol A, as described in figure 2, B the random variable describing the bits output by $\mathcal{B}$, and CHSH the event that the boxes' outputs are accepted in the final test described in the protocol. Then for all large enough n at least one of the following holds*:

— *either* $H_\infty^\varepsilon(B|CHSH) \geq n$,
— *or* $\Pr(CHSH) \leq \varepsilon$.

Furthermore, inputs in protocol A can be generated using $O(\log n \log(1/\varepsilon))$ bits of randomness, and it makes $O(n(\log n + \log(1/\varepsilon)))$ uses of the boxes. Finally, there exist honest, quantum-mechanical boxes that, when used in protocol A, are such that $\Pr(CHSH) \geq 1 - \varepsilon$.

Theorem 4.1 asserts that, given any pair $(\mathcal{A}, \mathcal{B})$ of non-signalling boxes, if the outputs of $\mathcal{B}$ do not contain large enough min-entropy (when its inputs are chosen as specified in the protocol) then the boxes will fail the CHSH constraints imposed in the protocol with a high probability.

The second part of the theorem follows from the definition of the protocol and the fact that there exists a good quantum strategy in the CHSH game, as was already demonstrated in §2. We prove the first, and main, part of theorem 4.1 by a reduction to the guessing game introduced in §3. Suppose that there existed a pair of boxes such that neither of the theorem's conclusions was satisfied. Recall that the protocol calls for a total of mk uses of the boxes, divided into m blocks of k pairs of identical inputs each. We show that, provided the CHSH constraints are satisfied in all blocks with non-negligible probability, there must exist a special block in which the boxes' outputs, conditioned on specific past values, have the properties required in lemma 3.1. This will in turn lead to a contradiction of the no-signalling assumption. The exact properties of the special block that we obtain are described in claim 4.2.

Modelling events in the protocol. Let $x = (x_i), y = (y_i), a = (a_i), b = (b_i) \in (\{0,1\}^k)^m$ denote the boxes' respective input and output strings in an execution of protocol A, as described in figure 2. Let X, Y, A, B be the corresponding random variables. The boxes' behaviour is characterized by a probability distribution $p_{AB|XY}(a, b|x, y)$. The iterative structure of the protocol implies that $p_{AB|XY}$ can be factored as follows:

$$p_{AB|XY}(a, b|x, y) = \prod_{i=1}^{mk} p_{A_iB_i|X_iY_iH_i}(a_i, b_i|x_i, y_i, h_i),$$

where $H_i = (A_{<i}, B_{<i}, X_{<i}, Y_{<i})$ and $h_i = (a_{<i}, b_{<i}, x_{<i}, y_{<i})$. We impose a single additional condition on $p_{AB|XY}$: that it obeys the no-signalling condition in every block, that is, for every $i \in [mk]$, a_i, x_i, y_i, y_i' and h_i, it holds that

$$\sum_{b_i} p_{A_iB_i|X_iY_iH_i}(a_i, b_i|x_i, y_i, h_i) = \sum_{b_i} p_{A_iB_i|X_iY_iH_i}(a_i, b_i|x_i, y_i', h_i),$$

and a symmetric condition holds when marginalizing over a_i. For $i \in [m]$, let CHSH_i be the event that $d_{\mathrm{H}}(A_i \oplus B_i, X_i \wedge Y_i) \leq 0.2$, and $\mathrm{CHSH} = \bigwedge_i \mathrm{CHSH}_i$. We will also use the shorthand $\mathrm{CHSH}_{<i} = \bigwedge_{j<i} \mathrm{CHSH}_j$. Finally, we let T_j be a random variable denoting the j-th Bell block, i.e. the j-th element of the set T chosen by the user in step 2 of protocol A.

Claim 4.2. Let n be an integer, and $2^{-n/10} < \varepsilon < \frac{1}{100}$. Suppose that (i) $H_\infty^\varepsilon(B|CHSH) \leq n$ and (ii) $\Pr(\mathrm{CHSH}) \geq \varepsilon$. Then for all large enough n there exists an index j_0 and a subset G of output strings satisfying $\Pr(B \in G) \geq \varepsilon^3$ such that the following hold:

— Conditioned on $\mathcal{B}$'s input in the j_0-th Bell block T_{j_0} being 0, its output in that block is essentially deterministic:

$$\forall b \in G, \quad \Pr(B_{T_{j_0}} = b_{T_{j_0}} | \mathrm{CHSH}_{<T_{j_0}}, B_{<T_{j_0}} = b_{<T_{j_0}}, Y_{T_{j_0}} = 0) \geq 0.95. \quad (4.1)$$

— The CHSH condition is satisfied with high probability in the j_0-th Bell block T_{j_0}:

$$\forall b \in G, \quad \Pr(\mathrm{CHSH}_{T_{j_0}} | \mathrm{CHSH}_{<T_{j_0}}, B_{<T_{j_0}} = b_{<T_{j_0}}) \geq 0.95. \tag{4.2}$$

The proof of claim 4.2 follows from an appropriate chained application of Baye's rule, and it is given in appendix A. In order to conclude the proof of theorem 4.1, it remains to show how the special block identified in claim 4.2 can lead to a successful strategy in the guessing game.

Consider the following strategy for Alice and Bob in the guessing game. In a preparatory phase (before Bob receives his secret bit y), Alice and Bob execute protocol A with the boxes $\mathcal{A}$ and $\mathcal{B}$, up to the T_{j_0}-th block (excluded). Bob communicates $\mathcal{B}$'s outputs up till that block to Alice. Together they check that the CHSH constraint is satisfied in all blocks preceding the T_{j_0}-th; if not, they abort. They also verify that Bob's outputs are the prefix of a string $b \in G$; if not, they abort. The guessing game can now start: Alice and Bob are separated and Bob is given his secret input y.

Given the conditioning that Alice and Bob have performed, once they are ready to start the game the boxes $\mathcal{A}$ and $\mathcal{B}$ satisfy both conditions of lemma 3.1. Indeed, under the input distribution specified in the protocol, inputs in a Bell block are chosen according to the uniform distribution. Because the block T_{j_0} identified in claim 4.2 is a Bell block by definition, condition 1 in lemma 3.1 holds with $\gamma = \frac{1}{20}$ as a consequence of item 1 in claim 4.2 and condition 2 in lemma 3.1 follows from item 2 in claim 4.2 with $\beta = \frac{1}{20}$. Because $\gamma/2 + 2\beta = 0.125 < \frac{1}{4}$, lemma 3.1 lets us conclude that the boxes $\mathcal{A}$ and $\mathcal{B}$ must be signalling in the T_{j_0}-th block, a contradiction. This finishes the proof of theorem 4.1.

U.V. is supported in part by NSF grant CCF-0905626, ARO grant W911NF-09-1-0440 and NIST award 60NANB10D262. T.V. is supported by the National Science Foundation under grant no. 0844626. Most of this work was completed while T.V. was at U.C. Berkeley.

Appendix A. Proof of claim 4.2

Before proving claim 4.2, we introduce the following well-known claim that will be useful in understanding what it means for a random variable to have small smooth min-entropy.

Claim A.1. Let $\alpha, \varepsilon > 0$ and X a random variable such that $H_\infty^\varepsilon(X) \leq \alpha$. Then there exists a set B such that $\Pr(X \in B) \geq \varepsilon$ and, for every $x \in B$, it holds that $\Pr(X = x) \geq 2^{-\alpha}$.

Proof. Let B be the set of x such that $\Pr(X = x) \geq 2^{-\alpha}$, and suppose $\Pr(X \in B) < \varepsilon$. Define Y so that $\Pr(Y = x) = \Pr(X = x)$ for every $x \notin B$, $\Pr(Y = x) = 0$ for every $x \in B$. In order to normalize Y, introduce new values z such that $\Pr(X = z) = 0$ and extend Y by defining $\Pr(Y = z) = 2^{-\alpha-1}$ until it is properly normalized. Then $\|Y - X\|_1 < \varepsilon$ and $H_\infty(Y) > \alpha$, contradicting the assumption on the smooth min-entropy of X. ∎

Proof of claim 4.2. As in protocol A, set $m = 400n$, $\Delta = 100\lceil \ln(1/\varepsilon) \rceil$ and $\ell = m/\Delta$. Let BAD be the set of strings $b \in (\{0,1\}^k)^m$ such that $\Pr(B = b | \mathrm{CHSH}) > 2^{-n}$. Assumption (i), together with claim A.1, shows that

$\Pr(\mathrm{BAD}|\mathrm{CHSH}) \geq \varepsilon$. Using (ii) and Baye's rule, we get that, for every $b = (b_1, \ldots, b_m) \in \mathrm{BAD}$,

$$\Pr(B = b, \mathrm{CHSH}) = \prod_{i=1}^{m} \Pr(B_i = b_i, \mathrm{CHSH}_i | \mathrm{CHSH}_{<i}, B_{<i} = b_{<i}) > 2^{-n}\varepsilon.$$

Taking logarithms on both sides,

$$\sum_{i=1}^{m} -\log \Pr(B_i = b_i, \mathrm{CHSH}_i | \mathrm{CHSH}_{<i}, B_{<i} = b_{<i}) < n + \log\left(\frac{1}{\varepsilon}\right) \leq \left(1 + \frac{1}{10}\right) n,$$

assuming as in the statement of the claim that ε is not too small. By an averaging argument, at least $\frac{9}{10}$ of all $i \in [m]$ are such that a fraction at least $\frac{6}{10}$ (in probability) of all $b \in \mathrm{BAD}$ are such that

$$\Pr(B_i = b_i | \mathrm{CHSH}_{<i}, B_{<i} = b_{<i}) \geq 2^{-(100/4)(1+1/10)(n/m)} \geq 2^{-28/400} \geq 0.95. \quad (\mathrm{A}\,1)$$

Let S be the set of $i \in [m]$ such that (A 1) holds for a fraction at least $\frac{6}{10}$ of $b \in \mathrm{BAD}$. S is a random variable of size $|S| \geq (\frac{9}{10})m$.

We apply the same reasoning once more, focusing on the CHSH constraint being satisfied in a Bell block. Let T be a random variable containing the indices of the blocks that have been designated as Bell blocks in the protocol. Let $N = |T \cap S|$. We may write N as the sum of Boolean random variables N_j, where $N_j = 1$ if and only if the j-th element of S falls in T. Because, for every i, the i-th block is chosen to be a Bell block independently with probability $1/\ell = \Delta/m$ (independently of past events such as $\mathrm{CHSH}_{<i}$ and $B_{<i} = b_{<i}$), the random variables N_j, for $j \leq |S|$, are independent. Recall that $|S| \geq 9m/10$, and by a Chernoff bound

$$\Pr\left(N_1 + \cdots + N_{9m/10} \geq \frac{9m}{10} \cdot \frac{1}{(10\ell)}\right) \geq 1 - \mathrm{e}^{-(9m/10\ell)(9/10)^2/2} \geq 1 - \mathrm{e}^{-\Delta/3} \geq 1 - \varepsilon^3,$$

given our choice of Δ. Let K denote the event that this bound holds: $\Pr(K) \geq 1 - \varepsilon^3$, and conditioned on K it holds that $N = |S \cap T| \geq 9m/(100\ell) \geq 9\Delta/100$. Starting from $\Pr(\mathrm{CHSH}|\mathrm{BAD}) \geq \varepsilon^2$, further conditioning on K gives

$$\Pr(\mathrm{CHSH}|\mathrm{BAD}, K) = \frac{\Pr(\mathrm{CHSH}, K|\mathrm{BAD})}{\Pr(K|\mathrm{BAD})} \geq \varepsilon^2 - \varepsilon^3 \geq \frac{\varepsilon^2}{2}.$$

Using Baye's rule as before, we then obtain

$$\sum_{i \in T \cap S} -\log \Pr(\mathrm{CHSH}_i | \mathrm{CHSH}_{<i}, \mathrm{BAD}_{<i}, K) \leq 2 \log\left(\frac{2}{\varepsilon}\right).$$

Using the lower bound on N, this implies that there exists an $i \in T \cap S$ such that

$$\Pr(\mathrm{CHSH}_i | \mathrm{CHSH}_{<i}, \mathrm{BAD}_{<i}, K) \geq \mathrm{e}^{2\log(2/\varepsilon)/N} \geq 0.956, \quad (\mathrm{A}\,2)$$

given the choice of Δ made in the claim. Given our assumption on ε, removing the conditioning on K in this equation at most decreases the lower bound to 0.95.

Let $i \in T \cap S$ be a Bell block for which (A 2) holds. By Markov's inequality, for a fraction at least $\frac{1}{2}$ of $b \in \mathrm{BAD}$, it holds that

$$\Pr(\mathrm{CHSH}_i | \mathrm{CHSH}_{<i}, B_{<i} = b_{<i}) \geq 0.9. \tag{A 3}$$

By the union bound, at iteration i (A 3) will hold simultaneously with (A 1) for a subset G of BAD of size at least

$$\Pr(G) = \Pr(G|\mathrm{BAD}) \Pr(\mathrm{BAD}) \geq (\tfrac{6}{10} - \tfrac{1}{2})\varepsilon^2 \geq \varepsilon^3$$

given our choice of parameters. Equation (A 3) implies (4.2) in the claim. To show that (A 1) implies (4.1), note that, in (A 1), the probability is taken for a random choice of inputs to the boxes in round i as specified in the protocol, while in (A 3), by definition of a Bell block, the probability is taken under the uniform distribution over inputs. We may not guarantee that (A 1) still holds under the uniform distribution, because in the protocol that distribution is chosen in round i only with probability $1/\ell$. However, because in the protocol Bob's input is a 0 with probability at least $\frac{1}{2}$ irrespective of the type of block, we may ensure that (A 1) holds conditioned on $Y_i = 0$, an event that is independent from both $\mathrm{CHSH}_{<i}$ and $B_{<i} = b_{<i}$ (for any $b_{<i}$). Hence, given our choice of Δ and the upper bound on ε, (A 1) implies (4.1) in the claim. ■

References

1 Campbell-Kelly, M. 1980 Programming the mark. I. Early programming activity at the university of Manchester. *Ann. Hist. Comput.* **2**, 136. (doi:10.1109/MAHC.1980.10018)

2 Taylor, G. & Cox, G. 2011 Behind intel's new random-number generator. *IEEE Spectr.* **September**, 32–35.

3 Downey, R. & Hirschfeldt, D. 2010 *Algorithmic randomness and complexity.* Berlin, Germany: Springer.

4 Goldreich, O. 2010 Pseudorandom generators: a primer. See http://www.wisdom.weizmann.ac.il/õded/prg-primer.html.

5 Santha, M. & Vazirani, U. V. 1984 Generating quasi-random sequences from slightly-random sources. In *Proc. of the 25th Annu. Symp. on Foundations of Computer Science, Singer Island, FL, 24–26 October 1984*, pp. 434–440. Washington, DC: IEEE Computer Society.

6 Zuckerman, D. 1990 General weak random sources. In *Proc. 31st Annu. Symp. on Foundations of Computer Science, St. Louis, MO, 22–24 October 1990*, pp. 534–543. Washington, DC: IEEE Computer Society.

7 Shaltiel, R. 2002 Recent developments in explicit constructions of extractors. *Bull. Eur. Assoc. Theor. Comput. Sci.* **77**, 67–95.

8 Stefanov, A., Gisin, N., Guinnard, O., Guinnard, L. & Zbinden, H. 2000 Optical quantum random number generator. *J. Mod. Opt.* **47**, 595–598. (http://arxiv.org/abs/qrant-ph/9907006)

9 Wang, P. X., Long, G. L. & Li, Y. S. 2006 Scheme for a quantum random number generator. *J. Appl. Phys.* **100**, 056 107. (doi:10.1063/1.2338830)

10 Svozil, K. 1990 The quantum coin toss-testing microphysical undecidability. *Phys. Lett. A* **143**, 433–437. (doi:10.1016/0375-9601(90)90408-G)

11 Fiorentino, M., Santori, C., Spillane, S. M., Beausoleil, R. G. & Munro, W. J. 2007 Secure self-calibrating quantum random-bit generator. *Phys. Rev. A.* **75**, 032 334. (doi:10.1103/PhysRevA.75.032334)

12 Calude, C., Dinneen, M., Dumitrescu, M. & Svozil, K. 2010 Experimental evidence of quantum randomness incomputability. *Phys. Rev. A.* **82**, 022102. (doi:10.1103/PhysRevA.82.022102)

13 Colbeck, R. 2009 *Quantum and relativistic protocols for secure multi-party computation.* Cambridge, UK: University of Cambridge.

14 Pironio, S. *et al.* 2009 Random numbers certified by Bell's theorem. *Nature* **464**, 1021–1024. (doi:10.1038/nature09008)
15 Einstein, A., Podolsky, P. & Rosen, N. 1935 Can quantum-mechanical description of physical reality be considered complete? *Phys. Rev.* **47**, 777–780. (doi:10.1103/PhysRev.47.777)
16 Bell, J. S. 1964 On the Einstein–Podolsky–Rosen paradox. *Physics* **1**, 195–200.
17 Bell, J. S. 1966 On the problem of hidden variables in quantum theory. *Rev. Mod. Phys.* **38**, 447–452. (doi:10.1103/RevModPhys.38.447)
18 Clauser, J. F., Horne, M. A., Shimony, A. & Holt, R. A. 1969 Proposed experiment to test local hidden-variable theories. *Phys. Rev. Lett.* **23**, 880–884. (doi:10.1103/PhysRevLett.23.880)
19 Colbeck, R. & Kent, A. 2011 Private randomness expansion with untrusted devices. *J. Phys. A Math. Theor.* **44**, 095 305. (doi:10.1088/1751-8113/44/9/095305)
20 Mayers, D. & Yao, A. 2004 Self testing quantum apparatus. *Quantum Inf. Comput.* **4**, 273–286. (http://arxiv.org/abs/grant.ph./0307205)
21 Barrett, J., Hardy, L. & Kent, A. 2005 No signaling and quantum key distribution. *Phys. Rev. Lett.* **95**, 010 503. (doi:10.1103/PhysRevLett.95.010503)
22 Fehr, S., Gelles, R. & Schaffner, C. 2011 Security and composability of randomness expansion from Bell inequalities. (http://arxiv.org/abs/1111.6052)
23 Bennett, C. & Brassard, G. 1984 Quantum cryptography: public key distribution and coin tossing. In *Proc. of IEEE Int. Conf. on Computers, Systems, and Signal Processing, Bangalore, India, 9–12 December 1984*, pp. 175–179. Washington, DC: IEEE Computer Society.
24 Pironio, S. & Massar, S. 2011 Security of practical private randomness generation. (http://arxiv.org/abs/1111.6056)
25 Terhal, B. 2004 Is entanglement monogamous? *IBM J. Res. Dev.* **48**, 71–78. (doi:10.1147/rd.481.0071)
26 Vazirani, U. V. & Vidick, T. 2012 Certifiable quantum dice—or, true random number generation secure against quantum adversaries. (http:arxiv.org/abs/1111.6054)
27 Knuth, A. D. Y. 1976 The complexity of nonuniform random number generation. In *Algorithms and complexity: new directions and recent results* (ed. J. F. Traub). New York, NY: Academic Press.
28 Guruswami, V., Umans, C. & Vadhan, S. 2009 Unbalanced expanders and randomness extractors from Parvaresh–Vardy codes. *J. ACM* **56**, 20:1–20:34.