Full randomness from arbitrarily deterministic events

Rodrigo Gallego[1,2], Lluis Masanes[1], Gonzalo De La Torre[1], Chirag Dhara[1],
Leandro Aolita[1,2], Antonio Acín[1,3]

[1]*ICFO-Institut de Ciències Fotòniques, Mediterranean Technology Park, 08860 Castelldefels (Barcelona), Spain*
[2]*Dahlem Center for Complex Quantum Systems, Freie Universität Berlin, Berlin, Germany*
[3]*ICREA-Institució Catalana de Recerca i Estudis Avançats, Lluís Companys 23, 08010 Barcelona, Spain*

Abstract

Do completely unpredictable events exist? Classical physics excludes fundamental randomness. While quantum theory makes probabilistic predictions, this does not imply that nature is random, as randomness should be certified without relying on the complete structure of the theory being used. Bell tests approach the question from this perspective. However, they require prior perfect randomness, falling into a circular reasoning. A Bell test that generates perfect random bits from bits possessing high -but less than perfect- randomness has recently been obtained. Yet, the main question remained open: does any initial randomness suffice to certify perfect randomness? We answer this question affirmatively. We provide a Bell test that uses arbitrarily imperfect random bits to produce bits that are, under the non-signalling principle assumption, perfectly random. This provides the first protocol attaining full randomness amplification. Our results have strong implications onto the debate of whether there exist events that are fully random.

Understanding whether nature is deterministically pre-determined or there are intrinsically random processes is a fundamental question that has attracted the interest of multiple thinkers, ranging from philosophers and mathematicians to physicists or neuroscientists. Nowadays this question is also important from a practical perspective, as random bits constitute a valuable resource for applications such as cryptographic protocols, gambling, or the numerical simulation of physical and biological systems.

Classical physics is a deterministic theory. Perfect knowledge of the positions and velocities of a system of classical particles at a given time, as well as of their interactions, allows one to predict their future (and also past) behavior with total certainty [1]. Thus, any randomness observed in classical systems is not intrinsic to the theory but just a manifestation of our imperfect description of the system.

The advent of quantum physics put into question this deterministic viewpoint, as there exist experimental situations for which quantum theory gives predictions only in probabilistic terms, even if one has a perfect description of the preparation and interactions of the system. A possible solution to this classically counterintuitive fact was proposed in the early days of quantum physics: Quantum mechanics had to be incomplete [2], and there should be a complete theory capable of providing deterministic predictions for all conceivable experiments. There would thus be no room for intrinsic randomness, and any apparent randomness would again be a consequence of our lack of control over hypothetical "hidden variables" not contemplated by the quantum formalism.

Bell's no-go theorem [3], however, implies that local hidden-variable theories are inconsistent with quantum mechanics. Therefore, none of these could ever render a deterministic completion to the quantum formalism. More precisely, all hidden-variable theories compatible with a local causal structure predict that any correlations among space-like separated events satisfy a series of inequalities, known as Bell inequalities. Bell inequalities, in turn, are violated by some correlations among quantum particles. This form of correlations defines the phenomenon of quantum non-locality.

Now, it turns out that quantum non-locality does not necessarily imply the existence of fully unpredictable processes in nature. The reasons behind this are subtle. First of all, unpredictable processes

could be certified only if the no-signaling principle holds. This states that no instantaneous communication is possible, which imposes in turn a local causal structure on events, as in Einstein's special relativity. In fact, Bohm's theory is both deterministic and able to reproduce all quantum predictions [4], but it is incompatible with no-signaling at the level of the hidden variables. Thus, we assume throughout the validity of the no-signaling principle. Yet, even within the no-signaling framework, it is still not possible to infer the existence of fully random processes only from the mere observation of non-local correlations. This is due to the fact that Bell tests require measurement settings chosen at random, but the actual randomness in such choices can never be certified. The extremal example is given when the settings are determined in advance. Then, any Bell violation can easily be explained in terms of deterministic models. As a matter of fact, super-deterministic models, which postulate that all phenomena in the universe, including our own mental processes, are fully pre-determined, are by definition impossible to rule out. These considerations imply that the strongest result on the existence of randomness one can hope for using quantum non-locality is stated by the following possibility: Given a source that produces an arbitrarily small but non-zero amount of randomness, can one still certify the existence of completely random processes?

In this work, we provide an affirmative answer to this question for a very general, and physically-meaningful, set of randomness sources. This includes subsets of the well known Santha-Vazirani sources [5] as particular cases. Besides the philosophical and physics-foundational implications, our results provide a protocol for full randomness amplification using quantum non-locality. Randomness amplification is an information-theoretic task whose goal is to use an input source of imperfectly random bits to produce perfect random bits. Santha and Vazirani proved that randomness amplification is impossible using classical resources [5]. This is in a sense intuitive, in view of the absence of any intrinsic randomness in classical physics. In the quantum regime, randomness amplification has been recently studied by Colbeck and Renner [6]. They proved how input bits with very high initial randomness can be mapped into arbitrarily pure random bits, and conjectured that randomness amplification should be possible for any initial randomness [6]. Our results also solve this conjecture, as we show that quantum non-locality can be exploited to attain *full randomness amplification*.

Before presenting our results, it is worth commenting on previous works on randomness in connection with quantum non-locality. In [7] it was shown how to bound the intrinsic randomness generated in a Bell test. These bounds can be used for device-independent randomness expansion, following a proposal by Colbeck [8], and to achieve a quadratic expansion of the amount of random bits (see [9, 10, 11, 12] for further works on device-independent randomness expansion). Note however that, in randomness expansion, one assumes instead, from the very beginning, the existence of an input seed of free random bits, and the main goal is to expand this into a larger sequence. The figure of merit there is the ratio between the length of the final and initial strings of free random bits. Finally, other recent works have analyzed how a lack of randomness in the measurement choices affects a Bell test [13, 14, 15] and the randomness generated in it [16].

Results

Definition of the scenario

From an information perspective, our goal is to construct a protocol for full randomness amplification based on quantum non-locality. In randomness amplification, one aims at producing arbitrarily free random bits from many uses of an input source $\mathcal{S}$ of imperfectly random bits.

A random bit b is said to be free if it is uncorrelated from any classical variables e generated outside the future light-cone of b (of course, the bit b can be arbitrarily correlated with any event inside its future light-cone). This requirement formalises the intuition that the only systems that may share some correlation with b are the ones that are influenced by b. Note also that this definition of randomness

Figure 1: **Local causal structure and randomness amplification**. A source $\mathcal{S}$ produces a sequence $x_1, x_2, \ldots x_j, \ldots$ of imperfect random bits. The goal of randomness amplification is to produce a new source $\mathcal{S}_f$ of perfect random bits, that is, to process the initial bits so as to get a final bit k fully uncorrelated (free) from all events outside the future light cone of all the bits x_j produced by the source. In other words k is free if it is uncorrelated from any event outside the lightcone shown in this figure. Any such event can be modeled by a measurement z, with an outcome e, on some physical system. This system may be under the control of an adversary Eve, interested in predicting the value of k.

is strictly stronger than the demand that b is uncorrelated with any classical variable generated in the past light-cone of the process. This is crucial if the variable e and b are generated by measuring on a correlated quantum system. In this case, even if both systems interacted somewhere in the past light-cone of b, the variable e is not produced until the measurement is performed, possibly outside both past and future light-cones. Furthermore, we say that a random bit is ϵ-free if any correlations with events outside its future light-cone are bounded by ϵ, as explained in what follows.

Source $\mathcal{S}$ produces a sequence of bits $x_1, x_2, \ldots x_j, \ldots$, with $x_j = 0$ or 1 for all j, see Fig. 1, which are ϵ-free. More precisely, each bit j contains some randomness, in the sense that the probability $P\left(x_j|\text{all other bits}, e\right)$ that it takes a given value x_j, conditioned on the values of all the other bits produced by $\mathcal{S}$, as well as the variable e, is such that

$$\epsilon \leq P\left(x_j|\text{all other bits}, e\right) \leq 1 - \epsilon \tag{1}$$

for all j, where $0 < \epsilon \leq 1/2$. Given our previous definition of ϵ-free bits, the variable e represents events outside the future light-cone of all the x_j's. Free random bits correspond to $\epsilon = \frac{1}{2}$; while deterministic ones to $\epsilon = 0$. More precisely, when $\epsilon = 0$ the bound (1) is trivial and no randomness can be certified. We refer to $\mathcal{S}$ as an ϵ-source, and to any bit satisfying (1) as an ϵ-free bit.

The aim of randomness amplification is to generate, from arbitrarily many uses of $\mathcal{S}$, a final source $\mathcal{S}_f$ of ϵ_f arbitrarily close to $1/2$. If this is possible, no cause e can be assigned to the bits produced by $\mathcal{S}_f$, which are then fully unpredictable. Note that, in our case, we require the final bits to be fully uncorrelated from e. Efficiency issues, such as the rate of uses of $\mathcal{S}$ required per final bit generated by $\mathcal{S}_f$ do not play any role in randomness amplification. The relevant figure of merit is just the quality, measured by ϵ_f, of the final bits. Thus, without loss of generality, we restrict our analysis to the problem of generating a single final free random bit k.

The randomness amplification protocols we consider exploit quantum non-locality. This idea was introduced in [6], where a protocol was presented in which the source $\mathcal{S}$ is used to choose the measurement settings by two distant observers, Alice and Bob, in a Bell test [17] involving two entangled

quantum particles. The measurement outcome obtained by one of the observers, say Alice, in one of the experimental runs (also chosen with $\mathcal{S}$) defines the output random bit. Colbeck and Renner proved how input bits with high randomness, of $0.442 < \epsilon \leq 0.5$, can be mapped into arbitrarily free random bits of $\epsilon_f \to 1/2$. In our case, the input ϵ-source $\mathcal{S}$ is used to choose the measurement settings in a multipartite Bell test involving a number of observers that depends both on the input ϵ and the target ϵ_f. After verifying that the expected Bell violation is obtained, the measurement outcomes are combined to define the final bit k. For pedagogical reasons, we adopt a cryptographic perspective and assume the worst-case scenario where all the devices we use may have been prepared by an adversary Eve equipped with arbitrary non-signaling resources, possibly even supra-quantum ones. In the preparation, Eve may have also had access to $\mathcal{S}$ and correlated the bits it produces with some physical system at her disposal, represented by a black box in Fig. 1. Without loss of generality, we can assume that Eve can reveal the value of e at any stage of the protocol by measuring this system. Full randomness amplification is then equivalent to proving that Eve's correlations with k can be made arbitrarily small.

An important comment is now in order that applies to all further discussion as well as the protocol subsequently presented. For convenience we represent (see Figs. 1 and 2) $\mathcal{S}$ as a single source generating all the inputs and delivering them among the separated boxes without violating the no-signalling principle. However, this is not the scenario in practice. Operationally, each user generates his input from a local source in his lab. However, all these sources can be arbitrarily correlated the sources of the other users with each other, without violating the bound on the correlations given by (1), and thus, can be seen as a single ϵ-source $\mathcal{S}$. With this understanding we proceed to discuss a single effective source in the rest of the text.

Partial randomness from GHZ-type paradoxes

Bell tests for which quantum correlations achieve the maximal non-signaling violation, also known as Greenberger-Horne-Zeilinger (GHZ)- type paradoxes [18], are necessary for full randomness amplification. This is due to the fact that unless the maximal non-signaling violation is attained, for sufficiently small ϵ, Eve may fake the observed correlations with classical deterministic resources. Nevertheless, GHZ-type paradoxes are not sufficient. In fact, given any function of the measurement outcomes, it is always possible to find non-signaling correlations that (i) maximally violate the 3-party GHZ paradox[18] but (ii) assign a deterministic value to that function of the measurement outcomes. This observation can be checked for all unbiased functions mapping $\{0,1\}^3$ to $\{0,1\}$ (there are $\binom{8}{4}$ of those) through a linear program analogous to the one used in the proof of the Lemma below. As a simple example, consider the particular function defined by the outcome bit of the first user. This can be fixed by using a tripartite no-signaling probability distribution consisting on a deterministic distribution for the first party, and a Popescu-Rohrlich (PR) box [19] for the second and third party.

For five parties though, the latter happens not to hold any longer. Consider now any correlations attaining the maximal violation of the five-party Mermin inequality [?]. In each run of this Bell test, measurements (inputs) $\mathbf{x} = (x_1, \ldots, x_5)$ on five distant black boxes generate 5 outcomes (outputs) $\mathbf{a} = (a_1, \ldots, a_5)$, distributed according to a non-signaling conditional probability distribution $P(\mathbf{a}|\mathbf{x})$, see Supplementary Note 1. Both inputs and outputs are bits, as they can take two possible values, $x_i, a_i \in \{0,1\}$ with $i = 1, \ldots, 5$.

The inequality can be written as

$$\sum_{\mathbf{a},\mathbf{x}} I(\mathbf{a},\mathbf{x}) P(\mathbf{a}|\mathbf{x}) \geq 6 \,, \tag{2}$$

with coefficients

$$I(\mathbf{a},\mathbf{x}) = (a_1 \oplus a_2 \oplus a_3 \oplus a_4 \oplus a_5)\, \delta_{\mathbf{x} \in \mathcal{X}_0} + (a_1 \oplus a_2 \oplus a_3 \oplus a_4 \oplus a_5 \oplus 1)\, \delta_{\mathbf{x} \in \mathcal{X}_1} \,, \tag{3}$$

where

$$\delta_{\mathbf{x}\in\mathcal{X}_s} = \left\{ \begin{array}{ll} 1 & \text{if } \mathbf{x} \in \mathcal{X}_s \\ 0 & \text{if } \mathbf{x} \notin \mathcal{X}_s \end{array} \right. ,$$

$$\mathcal{X}_0 = \left\{ \mathbf{x} \ \Big| \ \sum_{i=1}^{5} x_i = 1 \right\} \cup \left\{ \mathbf{x} \ \Big| \ \sum_{i=1}^{5} x_i = 5 \right\}, \tag{4}$$

and

$$\mathcal{X}_1 = \left\{ \mathbf{x} \ \Big| \ \sum_{i=1}^{5} x_i = 3 \right\}. \tag{5}$$

That is, only half of all possible combinations of inputs, namely those in $\mathcal{X} = \mathcal{X}_0 \cup \mathcal{X}_1$, appear in the Bell inequality. This inequality may be thought of as a non-local game in which the parties are required to minimize the parity of their outputs when the sum of their inputs is 1 or 5 while minimizing the inverse parity of the outputs when their inputs sum to 3. It turns out that the minimum achievable with classical strategies is 6.

The maximal, non-signaling and algebraic, violation of the inequality corresponds to the situation in which the left-hand side of (2) is zero. The key property of inequality (2) is that its maximal violation can be attained by quantum correlations and furthermore, one can construct a function of the outcomes that is not completely determined. Take the bit corresponding to the majority-vote function of the outcomes of any subset of three out of the five observers, say the first three. This function is equal to zero if at least two of the three bits are equal to zero, and equal to one otherwise. We show that Eve's predictability on this bit is at most 3/4. We state this result in the following Lemma:

Lemma. Let a five-party non-signaling conditional probability distribution $P(\mathbf{a}|\mathbf{x})$ in which inputs $\mathbf{x} = (x_1, \ldots, x_5)$ and outputs $\mathbf{a} = (a_1, \ldots, a_5)$ are bits. Consider the bit $\mathrm{maj}(\mathbf{a}) \in \{0, 1\}$ defined by the majority-vote function of any subset consisting of three of the five measurement outcomes, say the first three, a_1, a_2 and a_3. Then, all non-signaling correlations attaining the maximal violation of the 5-party Mermin inequality are such that the probability that $\mathrm{maj}(\mathbf{a})$ takes a given value, say 0, is bounded by

$$1/4 \leq P\left(\mathrm{maj}(\mathbf{a}) = 0\right) \leq 3/4. \tag{6}$$

Proof. This result was obtained by solving a linear program. Therefore, the proof is numeric, but exact. Formally, let $P(\mathbf{a}|\mathbf{x})$ be a 5-partite non-signaling probability distribution. For $\mathbf{x} = \mathbf{x}_0 \in \mathcal{X}$, we performed the maximization,

$$P_{max} = \max_{P} P(\mathrm{maj}(\mathbf{a}) = 0|\mathbf{x}_0)$$

$$\text{subject to} \tag{7}$$

$$I(\mathbf{a}, \mathbf{x}) \cdot P(\mathbf{a}|\mathbf{x}) = 0$$

which yields the value $P_{max} = 3/4$. Since the same result holds for $P(\mathrm{maj}(\mathbf{a}) = 1|\mathbf{x}_0)$, we get the bound $1/4 \leq P(\mathrm{maj}(\mathbf{a}) = 0) \leq 3/4$.

$\square$

As a further remark, note that a lower bound to P_{max} can easily be obtained by noticing that one can construct conditional probability distributions $P(\mathbf{a}|\mathbf{x})$ that maximally violate 5-partite Mermin inequality (2) for which at most one of the output bits (say a_1) is deterministically fixed to either 0 or 1. If the other two output bits (a_2, a_3) were to be completely random, the majority-vote of the three of them $\mathrm{maj}(a_1, a_2, a_3)$ could be guessed with a probability of 3/4. Our numerical results say that this turns out to be an optimal strategy.

The previous lemma strongly suggests that, given an ϵ-source with any $0 < \epsilon \leq 1/2$ and quantum five-party non-local resources, it should be possible to design a protocol to obtain an ϵ_i-source of $\epsilon_i =$

1/4. We do not explore this possibility here, but rather use the partial unpredictability in the five-party Mermin Bell test as building block of our protocol for full randomness amplification. To complete it, we must equip it with two essential components: (*i*) an *estimation procedure* that verifies that the untrusted devices do yield the required Bell violation; and (*ii*) a *distillation procedure* that, from sufficiently many ϵ_i-bits generated in the 5-party Bell experiment, distills a single final ϵ_f-source of $\epsilon_f \rightarrow 1/2$. Towards these ends, we consider a more complex Bell test involving N groups of five observers (quintuplets) each.

A protocol for full randomness amplification

Our protocol for randomness amplification uses as resources the ϵ-source $\mathcal{S}$ and $5N$ quantum systems. Each of the quantum systems is abstractly modeled by a black box with binary input x and output a. The protocol processes classically the bits generated by $\mathcal{S}$ and by the quantum boxes. When the protocol is not aborted it produces a bit k. The steps of the protocol are described in what follows (see also Fig. 2).

1. $\mathcal{S}$ is used to generate N quintuple-bits $\mathbf{x}_1, \dots \mathbf{x}_N$, which constitute the inputs for the $5N$ boxes and are distributed among them without violating no-signaling. The boxes then provide N output quintuple-bits $\mathbf{a}_1, \dots \mathbf{a}_N$.

2. The quintuplets such that $\mathbf{x} \notin \mathcal{X}$ are discarded. The protocol is aborted if the number of remaining quintuplets is less than $N/3$. (Note that the constant factor $1/3$ is arbitrary. In fact, it is enough to demand that the number of remaining quintuplets is larger than N/c, with $c > 1$. See the Supplementary Note 2.)

3. The quintuplets left after step 2 are organized in N_b blocks each one having N_d quintuplets. The number N_b of blocks is chosen to be a power of 2. For the sake of simplicity, we relabel the index running over the remaining quintuplets, namely $\mathbf{x}_1, \dots \mathbf{x}_{N_b N_d}$ and outputs $\mathbf{a}_1, \dots \mathbf{a}_{N_b N_d}$. The input and output of the j-th block are defined as $y_j = (\mathbf{x}_{(j-1)N_d+1}, \dots \mathbf{x}_{(j-1)N_d+N_d})$ and $b_j = (\mathbf{a}_{(j-1)N_d+1}, \dots \mathbf{a}_{(j-1)N_d+N_d})$ respectively, with $j \in \{1, \dots, N_b\}$. The random variable $l \in \{1, \dots N_b\}$ is generated by using $\log_2 N_b$ further bits from $\mathcal{S}$. The value of l specifies which block (b_l, y_l) is chosen to generate k, i.e. the distilling block. We define $(\tilde{b}, \tilde{y}) = (b_l, y_l)$. The other $N_b - 1$ blocks are used to check the Bell violation.

4. The function
$$r[b, y] = \begin{cases} 1 & \text{if } I(\mathbf{a}_1, \mathbf{x}_1) = \cdots = I(\mathbf{a}_{N_d}, \mathbf{x}_{N_d}) = 0 \\ 0 & \text{otherwise} \end{cases} \tag{8}$$
tells whether block (b, y) features the right correlations ($r = 1$) or the wrong ones ($r = 0$), in the sense of being compatible with the maximal violation of inequality (2). This function is computed for all blocks but the distilling one. The protocols is aborted unless all of them give the right correlations,
$$g = \prod_{j=1, j \neq l}^{N_b} r[b_j, y_j] = \begin{cases} 1 & \text{not abort} \\ 0 & \text{abort} \end{cases}. \tag{9}$$
Note that the abort/no-abort decision is independent of whether the distilling block l is right or wrong.

5. If the protocol is not aborted then k is assigned a bit generated from $b_l = (\mathbf{a}_1, \dots \mathbf{a}_{N_d})$ as
$$k = f(\text{maj}(\mathbf{a}_1), \dots \text{maj}(\mathbf{a}_{N_d})). \tag{10}$$
Here $f : \{0, 1\}^{N_d} \rightarrow \{0, 1\}$ is a function whose existence is proven in the Supplementary Note 2, while $\text{maj}(\mathbf{a}_i) \in \{0, 1\}$ is the majority-vote among the three first bits of the quintuple string $\mathbf{a}_i$.
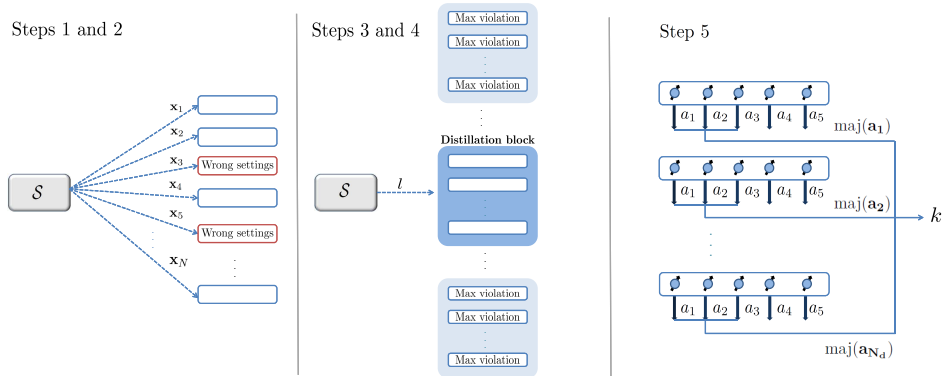
Figure 2: **Protocol for full randomness amplification based on quantum non-locality**. In the first two steps, all N quintuplets measure their devices, where the choice of measurement is done using the ϵ-source $\mathcal{S}$. Although it is illustrated here as a single source for convenience we recall that it represents the collection of sources that each space-like separated party locally possesses with all their outputs being correlated to form an ϵ-source. The quintuplets whose settings happen not to take place in the five-party Mermin inequality are discarded (in red). In steps 3 and 4, the remaining quintuplets are grouped into blocks. One of the blocks is chosen as the distillation block, using again $\mathcal{S}$, while the others are used to check the Bell violation. In the fifth step, the random bit k is extracted from the distillation block.

At the end of the protocol, the bit k is potentially correlated with the settings of the distilling block $\tilde{y} = y_l$, the bit g defined in (9), and the information

$$t = [l, (b_1, y_1), \dots (b_{l-1}, y_{l-1}), (b_{l+1}, y_{l+1}), \dots (b_{N_b}, y_{N_b})].$$

Additionally, an eavesdropper Eve might have access to a physical system correlated with k, which she can measure at any stage of the protocol. This system is not necessarily classical nor quantum, the only assumption about it is that measuring it does not produce instantaneous signaling anywhere else. The measurements that Eve can perform on her system are labeled by z, and the corresponding outcomes by e. In summary, after performing the protocol all the relevant information is $k, \tilde{y}, t, g, e, z$, with statistics described by an unknown conditional probability distribution $P(k, \tilde{y}, t, g, e|z)$. When the protocol is aborted ($g = 0$) there is no value for k. Therefore, in order to have a well defined distribution $P(k, \tilde{y}, t, g, e|z)$ in all cases, we set $k = 0$ when $g = 0$—that is $P(k, \tilde{y}, t, g = 0, e|z) = \delta_k^0 P(\tilde{y}, t, g = 0, e|z)$, where δ_k^0 is the Kronecker tensor.

To assess the quality of our protocol for full randomness amplification we compare it with an ideal protocol having the same marginal for the variables $\tilde{y}, t, g$ and the physical system described by e, z. That is, the global distribution of the ideal protocol is

$$P_{\text{ideal}}(k, \tilde{y}, t, g, e|z) = \begin{cases} \frac{1}{2} P(\tilde{y}, t, g, e|z) & \text{if } g = 1 \\ \delta_k^0 P(\tilde{y}, t, g, e|z) & \text{if } g = 0 \end{cases}, \tag{11}$$

7

where $P(\tilde{y}, t, g, e|z)$ is the marginal of the distribution $P(k, \tilde{y}, t, g, e|z)$ generated by the real protocol. Note that, consistently, in the ideal distribution we also set $k = 0$ when $g = 0$.

Our goal is that the statistics of the real protocol P is indistinguishable from the ideal statistics P_{ideal}. We consider the optimal strategy to discriminate between P and P_{ideal}, which obviously involves having access to all possible information $k, \tilde{y}, t, g$ and the physical system e, z. As shown in [21], the optimal probability for correctly guessing between these two distributions is

$$P(\text{guess}) = \frac{1}{2} + \frac{1}{4} \sum_{k, \tilde{y}, t, g} \max_z \sum_e \left| P(k, \tilde{y}, t, g, e|z) - P_{\text{ideal}}(k, \tilde{y}, t, g, e|z) \right| . \tag{12}$$

Note that the second term can be understood as (one half of) the variational distance between P and P_{ideal} generalized to the case when the distributions are conditioned on an input z. The following theorem is proven in the Supplementary Note 2.

Theorem. Let $P(k, \tilde{y}, t, g, e|z)$ be the probability distribution of the variables generated during the protocol and the adversary's physical system e, z; and let $P_{\text{ideal}}(k, \tilde{y}, t, g, e|z)$ be the corresponding ideal distribution (11). The optimal probability of correctly guessing between the distributions P and P_{ideal} satisfies

$$P(\text{guess}) \leq \frac{1}{2} + \frac{3\sqrt{N_d}}{2} \left[\alpha^{N_d} + 2 N_b^{\log_2(1-\epsilon)} \left(32\beta\epsilon^{-5} \right)^{N_d} \right] , \tag{13}$$

where the real numbers α, β fulfill $0 < \alpha < 1 < \beta$.

Now, the right-hand side of (13) can be made arbitrary close to $1/2$, for instance by setting $N_b = \left(32\,\beta\,\epsilon^{-5} \right)^{2N_d/|\log_2(1-\epsilon)|}$ and increasing N_d subject to the condition $N_d N_b > N/3$. [Note that $\log_2(1-\epsilon) < 0$.] In the limit of large N_d the probability $P(\text{guess})$ tends to $1/2$, which implies that the optimal strategy is as good as tossing a coin. In this case, the performance of the protocol is indistinguishable from that of an ideal one. This is known as "universally-composable security", and accounts for the strongest notion of cryptographic security (see [22] and [21]).

Let us discuss the implications and limitations of our result. Note first that step 2 in the protocol involves a possible abortion (a similar step can also be found in Ref. [6]). Hence, only those ϵ-sources with a non-negligible probability of passing step 2 can be amplified by our protocol. The abortion step can be relaxed by choosing a larger value of the constant c used for rejection. Yet, in principle, it could possibly exclude some of the ϵ-sources defined in (1). Notice, however, that demanding that step 2 is satisfied with non-negligible probability is just a restriction on the statistics seen by the honest parties $P(x_1, \ldots, x_n)$ and does not imply any restriction on the value of ϵ in $\epsilon \leq P(x_1, \ldots, x_n|e) \leq (1-\epsilon)$, which can be arbitrarily small. Also, we identify at least two reasons why sources that fulfil step 2 with high probability are the most natural in the context of randomness amplification. First, from a cryptographic perspective, if the observed sequence $x_1, \ldots, x_n$ does not fulfill step 2, then the honest parties will abort any protocol, regardless of whether a condition similar to step 2 is included. The reason is that such sequence would be extremely atypical in a fair source $P(x_1, \ldots, x_n) = 1/2^n$ and thus the honest players will conclude that the source is intervened by a malicious party or seriously damaged. Moreover, as discussed in the Supplementary Note 3, imposing that the source has unbiased statistics from the honest parties' point of view does not imply any restriction on Eve's predictability. Second, from a more fundamental viewpoint, the question of whether truly random events exist in Nature is interesting since the observable statistics of many physical processes look random, that is, they are such that $P(x_1, \ldots, x_n) = 1/2^n$. If every process in nature was such the observable statistics does not fulfill step 2, the problem of whether truly random processes exist would hardly have been considered relevant. Finally, note that possible sources outside this subclass do not compromise the security of the protocol, only its probability of being successfully implemented.

8

Under the conditions demanded in step 2, our protocol actually goes through for sources more general than those in (1). These are defined by the following restrictions,

$$G(n, \epsilon) < P(x_1, \ldots, x_n|e) \leq F(n, \epsilon), \tag{14}$$

for any pair of functions $G(n, \epsilon), F(n, \epsilon)$ defining the lower and upper bounds to Eve's control on the bias of each bit fulfilling the conditions $G(n, \epsilon) > 0$ and $\lim_{n \to \infty} F(n, \epsilon) = 0$. In fact, this condition is sufficient for our amplification protocol to succeed, see also the Supplementary Note 2 and 3.

To complete the argument we must mention that, according to quantum mechanics, given a source that passes step 2, we can in principle implement the protocol with success probability equal to one, $P(g = 1) = 1$. It can be immediately verified that the qubit measurements X or Y on the quantum state $|\Psi\rangle = \frac{1}{\sqrt{2}}(|00000\rangle + |11111\rangle)$, with $|0\rangle$ and $|1\rangle$ the eigenstates of Z, yield correlations that maximally violate the five-partite Mermin inequality in question. (In a realistic scenario the success probability $P(g = 1)$ might be lower than one, but our Theorem warrants that the protocol is still secure.)

We can now state the main result of our work.

Full randomness amplification: A perfect free random bit can be obtained from sources of arbitrarily weak randomness using non-local quantum correlations.

We would like to conclude by explaining the main intuitions behind the proof of the previous theorem. As mentioned, the protocol builds on the 5-party Mermin inequality because it is the simplest GHZ paradox allowing some randomness certification. The estimation part, given by step 4, is rather standard and inspired by estimation techniques introduced in [23], which were also used in [6] in the context of randomness amplification. The most subtle part is the distillation of the final bit in step 5. Naively, and leaving aside estimation issues, one could argue that it is nothing but a classical processing by means of the function f of the imperfect random bits obtained via the N_d quintuplets. But this seems in contradiction with the result by Santha and Vazirani proving that it is impossible to extract by classical means a perfect free random bit from imperfect ones [5]. This intuition is however misleading. Indeed, the Bell certification allows applying techniques similar to those obtained in Ref. [21] in the context of privacy amplification against non-signaling eavesdroppers. There, it was shown how to amplify the privacy, that is the unpredictability, of one of the measurement outcomes of bipartite correlations violating a Bell inequality. The key point is that the amplification, or distillation, is attained in a *deterministic* manner. That is, contrary to standard approaches, the privacy amplification process described in [21] does not consume any randomness. Clearly, these deterministic techniques are extremely convenient for our randomness amplification scenario. In fact, the distillation part in our protocol can be seen as the translation of the privacy amplification techniques of Ref. [21] to our more complex scenario, involving now 5-party non-local correlations and a function of three of the measurement outcomes.

Discussion

We have presented a protocol that, using quantum non-local resources, attains *full randomness amplification*, a task known to be impossible classically. As our goal was to prove full randomness amplification, our analysis focuses on the noise-free case. In fact, the noisy case only makes sense if one does not aim at perfect random bits and bounds the amount of randomness in the final bit. Then, it should be possible to adapt our protocol in order to get a bound on the noise it tolerates. Other open questions that our results offer as challenges consist of extending randomness amplification to other randomness sources, studying randomness amplification against quantum eavesdroppers, or the search of protocols in the bipartite scenario.

From a more fundamental perspective, our results imply that there exist experiments whose outcomes are fully unpredictable. The only two assumptions for this conclusion are the existence of events

with an arbitrarily small but non-zero amount of randomness that pass step 2 of our protocol and the validity of the no-signaling principle. Dropping the first assumption would lead to super-determinism, or to accept that the only source of randomness in nature are those that do not pass step 2 of our protocol, and in particular, that do not look unbiased. On the other hand, dropping the second assumption would imply abandoning a local causal structure for events in space-time. However, this is one of the most fundamental notions of special relativity.

References

[1] Laplace, P. S. A philosophical essay on probabilities. (1840).

[2] Einstein, A., Podolsky, B. & Rosen, N. Can Quantum-Mechanical Description of Physical Reality Be Considered Complete? *Phys. Rev.* **47**, 777-780 (1935).

[3] Bell, J. On the Einstein Podolsky Rosen Paradox. *Physics* **1**, 195-200 (1964).

[4] Bohm, D. A Suggested Interpretation of the Quantum Theory in Terms of "Hidden" Variables. I. *Phys. Rev.* **85**, 166–179 (1952).

[5] Santha M. & Vazirani U. V. Generating Quasi-random Sequences from Semi-random Sources. *J. Comput. Syst. Sci.* **33**, 75-87 (1986).

[6] Colbeck, R. & Renner, R. Free randomness can be amplified, *Nature Phys.* **8**, 450–454 (2012).

[7] Pironio S. *et al.* Random numbers certified by Bell's theorem. Nature **464**, 1021-1024 (2010).

[8] Colbeck R. Quantum and Relativistic Protocols for Secure Multi-Party Computation, *PhD Thesis*, Univ. of Cambridge (2007).

[9] Acín A., Massar S. & Pironio, S. Randomness versus Nonlocality and Entanglement. *Phys. Rev. Lett.* **108**, 100402 (2012).

[10] Pironio, S. & Massar, S. Security of practical private randomness generation. *Phys. Rev. A* **87**, 012336 (2013).

[11] Fehr, S., Gelles, R. & Schaffner, C. Security and Composability of Randomness Expansion from Bell Inequalities. *arXiv:1111.6052* (2011).

[12] Vazirani, U. & Vidick, T. Certifiable Quantum Dice: or, true random number generation secure against quantum adversaries. *Proceedings of the ACM Symposium on the Theory of Computing* (2012).

[13] Kofler, J., Paterek, T. & Brukner, C. Experimenter's freedom in Bell's theorem and quantum cryptography. *Phys. Rev. A* **73**, 022104 (2006).

[14] Barrett J. & Gisin N. How much measurement independence is needed in order to demonstrate nonlocality? *Phys. Rev. Lett* **106**, 100406 (2011).

[15] Hall, M. J. W. Local Deterministic Model of Singlet State Correlations Based on Relaxing Measurement Independence. *Phys. Rev. Lett.* **105**, 250404 (2010).

[16] Koh, D. E. *et al.* The effects of reduced "free will" on Bell-based randomness expansion. *Phys. Rev. Lett.* **109**, 160404 (2012).

[17] Braunstein, S. L. & Caves, C. M. Wringing out better Bell inequalities. *Annals of Physics* **202**, 22 (1990).

[18] Greenberger, D. M., Horne, M. A. & Zeilinger, A. Bell's Theorem, Quantum Theory, and Conceptions of the Universe. Kluwer, Dordrecht (1989).

[19] Popescu, S. & Rohrlich, D. Quantum nonlocality as an axiom. *Foundations of Physics* **24**, 379-385 (1994).

[20] Mermin, N. D. Simple unified form for the major no-hidden-variables theorems. *Phys. Rev. Lett.* **65**, 3373-3376 (1990).

[21] Masanes, L. Universally Composable Privacy Amplification from Causality Constraints. *Phys. Rev. Lett.* **102**, 140501 (2009).

[22] Canetti, R., Universally composable security: a new paradigm for cryptographic protocols. *Proc. 42nd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 136-145 (2001).

[23] Barrett, J., Hardy, L. & Kent, A. No Signaling and Quantum Key Distribution. *Phys. Rev. Lett.* **95**, 010503 (2005).

Acknowledgements

Author Contributions

All authors contributed extensively to the work presented in this paper.

Additional Information

Competing financial interests: The authors declare no competing financial interests.