LETTER

https://doi.org/10.1038/s41586-018-0019-0

Experimentally generated randomness certified by the impossibility of superluminal signals

Peter Bierhorst[1,2]*, Emanuel Knill[1,3], Scott Glancy[1], Yanbao Zhang[1,8], Alan Mink[4,5], Stephen Jordan[4], Andrea Rommal[6], Yi-Kai Liu[4], Bradley Christensen[7], Sae Woo Nam[1], Martin J. Stevens[1] & Lynden K. Shalm[1,2]

From dice to modern electronic circuits, there have been many attempts to build better devices to generate random numbers. Randomness is fundamental to security and cryptographic systems and to safeguarding privacy. A key challenge with random-number generators is that it is hard to ensure that their outputs are unpredictable[1–3]. For a random-number generator based on a physical process, such as a noisy classical system or an elementary quantum measurement, a detailed model that describes the underlying physics is necessary to assert unpredictability. Imperfections in the model compromise the integrity of the device. However, it is possible to exploit the phenomenon of quantum nonlocality with a loophole-free Bell test to build a random-number generator that can produce output that is unpredictable to any adversary that is limited only by general physical principles, such as special relativity[1–11]. With recent technological developments, it is now possible to carry out such a loophole-free Bell test[12–14,22]. Here we present certified randomness obtained from a photonic Bell experiment and extract 1,024 random bits that are uniformly distributed to within 10^{-12}. These random bits could not have been predicted according to any physical theory that prohibits faster-than-light (superluminal) signalling and that allows independent measurement choices. To certify and quantify the randomness, we describe a protocol that is optimized for devices that are characterized by a low per-trial violation of Bell inequalities. Future random-number generators based on loophole-free Bell tests may have a role in increasing the security and trust of our cryptographic systems and infrastructure.

The search for certifiably unpredictable random-number generators is motivated by applications, such as secure communication, for which the predictability of pseudorandom strings makes them unsuitable. Private randomness is required to initiate and authenticate virtually every secure communication[15], and public randomness from randomness beacons can be used for public certification and resource distribution in many settings[16]. To certify randomness, we can perform an experiment known as a Bell test[17]; in its simplest form, the Bell test involves performing measurements on an entangled system with components located in two physically separated measurement stations, where at each station a choice is made between one of two types of measurement. After multiple experimental trials with varying measurement choices, if the measurement data violate conditions known as 'Bell inequalities', then the data are certified to contain randomness under weak assumptions.

Our randomness generation uses a 'loophole-free' Bell test, which is characterized by high detection efficiency and space-like separation of the measurement stations during each experimental trial. The bits are unpredictable assuming (1) that the choices of measurement setting are independent of the experimental devices and of pre-existing classical information about them and (2) that, in each experimental trial, the measurement outcomes at each station are independent of the settings at the other station. The first assumption is ultimately untestable, but the premise that it is possible to choose measurement settings independently of a system being measured is often tacitly invoked in the interpretation of many scientific experiments and laws of physics[18]. The second assumption can be violated only if signals can be sent faster than the speed of light, given our trust that the space-like separation of the relevant events in the experiment is accurately verified by the timing electronics and that the results are final when recorded. We also trust that the classical computing equipment used to process the data operates according to specification.

Under the above assumptions, the output randomness is certified to be unpredictable with respect to a real or hypothetical actor 'Eve', who is in possession of the pre-existing classical information, is physically isolated from the devices while they are under our control and is without access to data produced during the protocol. The bits remain unpredictable to Eve if she learns the settings at any time after her last interaction with the devices. If the devices are trusted, which is reasonable if we built them, then this final interaction may be well before the start of the protocol, in which case the settings can come from public randomness[2,10]. In particular, an existing public randomness source can be used, such as the National Institute of Standards and Technology (NIST) random beacon[16], to generate much-needed private randomness as output. Because the assumptions do not constrain the specific physical realization of the devices and do not require specific states or measurements, they implement a 'device-independent' framework[4,19,20], which allows an individual user to assure security with minimal assumptions about the devices.

Compared to other implementations of random-number generations that invoke device-independence[5,21], our implementation is notable because it enforces space-like separation between measurement stations. Bell tests that achieve space-like separation without other experimental loopholes have been performed only recently[12–14,22]. It can be argued that interaction between spatially (if not space-like) separated measuring stations can be assumed to be negligible. However, any shielding between the stations is necessarily incomplete; for example, there must be an open quantum channel to establish entanglement. Mundane physical effects, such as accidentally scattered photons, can allow predictable systems to appear to violate Bell inequalities when shielding is incomplete. Relying instead on the impossibility of faster-than-light communication provides stronger assurance of the unpredictability of the randomness.

We generated randomness using an improved version of a recently reported[13] loophole-free Bell test (which was subsequently used elsewhere[23]). We collected five datasets, with the best-performing one yielding 1,024 random bits that are uniformly distributed to within 10^{-12}, as measured by the total variation distance (see below). We also obtained 256 random bits from the main dataset analysed previously[13], albeit uniform only to within 0.02; see Supplementary Information section 6. The experiment, illustrated in

[1]National Institute of Standards and Technology, Boulder, CO, USA. [2]Department of Physics, University of Colorado, Boulder, CO, USA. [3]Center for Theory of Quantum Matter, University of Colorado, Boulder, CO, USA. [4]National Institute of Standards and Technology, Gaithersburg, MD, USA. [5]Theiss Research, La Jolla, CA, USA. [6]Muhlenberg College, Allentown, PA, USA. [7]Department of Physics, University of Wisconsin, Madison, WI, USA. [8]Present address: NTT Basic Research Laboratories and NTT Research Center for Theoretical Quantum Physics, NTT Corporation, Atsugi, Japan. *e-mail: peter.bierhorst@nist.gov

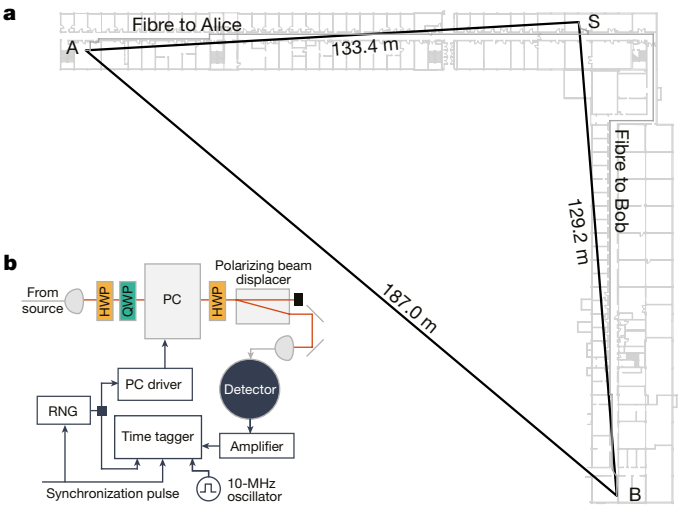

Fig. 1 | Diagram of the experiment. a, **b**, The relative locations of the source (S), Alice (A) and Bob (B) are depicted in **a**. In each trial, the source laboratory produces a pair of photons in a non-maximally polarization-entangled state. One photon is sent to Alice's laboratory while the other is sent to Bob's laboratory to be measured, as shown in **b**. Alice and Bob both use a fast Pockels cell (PC), two half-wave plates (HWPs), a quarter-wave plate (QWP) and a polarizing beam displacer to switch between their respective polarization measurements. A pseudorandom-number generator (RNG) governs the choice of each measurement setting for each trial. After passing through the polarization optics, the photons are sent to a superconducting nanowire detector. The signals from the detector are amplified and sent to a time tagger, where their arrival times are recorded and the measurement outcome is fixed. Alice's measurement outcome is space-like separated from the triggering of Bob's Pockels cell and vice versa.

Fig. 1, consisted of a source of entangled photons and two measurement stations, named 'Alice' and 'Bob'. During an experimental trial, at each station a random choice was made between two measurement settings, labelled 0 and 1, after which a measurement outcome of detection (+) or non-detection (0) was recorded. Each station's implementation of the measurement setting was space-like separated from the other station's measurement event, and no post-selection was used in collecting the data; see Methods for details. For trial i, we model Alice's settings choices with the random variable X_i and Bob's with Y_i, both of which take values in the set {0, 1}. Alice's and Bob's measurement-outcome random variables are A_i and B_i, respectively, both of which take values in the set {+, 0}. When referring to a generic single trial, we omit the i indices. With this notation, a general Bell inequality for our scenario can be expressed in the form[24]

$$\sum_{abxy} s_{xy}^{ab}\mathbb{P}(A=a,B=b \mid X=x,Y=y) \leq \beta \tag{1}$$

where s_{xy}^{ab} are fixed real coefficients indexed by a, b, x and y, which range over all possible values of A, B, X and Y, and $\mathbb{P}$ denotes probability. The upper bound β is required to be satisfied whenever the settings-conditional outcome probabilities are induced by a model that satisfies 'local realism'. Local-realist distributions, which cannot be certified to contain randomness, are those for which $\mathbb{P}(A=a,B=b \mid X=x,Y=y)$ is of the form $\sum_\lambda \mathbb{P}(A=a \mid X=x,\Lambda=\lambda)\mathbb{P}(B=b \mid Y=y,\Lambda=\lambda)\mathbb{P}(\Lambda=\lambda)$ for a random variable Λ that represents local hidden variables. The Bell inequality is non-trivial if there exists a quantum-realizable distribution that can violate the bound β.

It has long been known that experimental violations of Bell inequalities such as equation (1) indicate the presence of randomness in data. To quantify randomness with respect to Eve, we represent Eve's initial classical information by a random variable E. We formalize the assumption that measurement settings can be generated independently of the system being measured and of Eve's information with the following condition:

$$\begin{aligned}\mathbb{P}(X_i=x,Y_i=y \mid E=e,\mathrm{past}_i) &= \mathbb{P}(X_i=x,Y_i=y)\\ &= \frac{1}{4} \quad \forall x,y,e\end{aligned} \tag{2}$$

where past_i represents events prior to the ith trial, specifically including the trial settings and outcomes for trials 1 to $i-1$. Our other assumption, that measurement outcomes are independent of remote measurement choices, is formalized as follows:

$$\begin{aligned}&\mathbb{P}(A_i=a \mid X_i=x,Y_i=y,E=e,\mathrm{past}_i)\\ &\quad=\mathbb{P}(A_i=a \mid X_i=x,E=e,\mathrm{past}_i) \quad \text{and}\\ &\mathbb{P}(B_i=b \mid X_i=x,Y_i=y,E=e,\mathrm{past}_i)\\ &\quad=\mathbb{P}(B_i=b \mid Y_i=y,E=e,\mathrm{past}_i) \quad \forall a,b,x,y,e\end{aligned} \tag{3}$$

These equations are commonly referred to as the 'non-signalling' assumptions, although they are often stated without the conditionals E and past_i. Our space-like separation of settings and remote measurements provide assurance that the experiment obeys equation (3). If we were to assume that the measured systems obey quantum physics, then stronger constraints are possible[25,26].

Given equations (2) and (3), our protocol produces random bits in two sequential parts. For the first part, 'entropy production', we implement n trials of the Bell test, from which we compute a statistic V that is related to a Bell inequality (equation (1)). V quantifies the Bell violation and determines whether or not the protocol passes or aborts. If the protocol passes, then we certify an amount of randomness in the outcome string whether or not Eve has access to the setting string. In the second part, 'extraction', we process the outcome string into a shorter string of bits, the distribution of which is close to uniform. We used our customized implementation of the Trevisan extractor[27] derived from the framework of Mauerer, Portmann and Scholz[28] and the associated open-source code. We call this the TMPS algorithm; see Supplementary Information section 4 for details.

We applied a new method of certifying the amount of randomness in Bell tests. Previous methods for related models with various sets of assumptions[2–8,29,30] are ineffective in our experimental regime (see Supplementary Information section 7), which is characterized by a small per-trial violation of Bell inequalities. Other recent works that explore ways of effectively certifying randomness from a wider range of experimental regimes assume that measured states are independent and identically distributed (i.i.d.) or that the regime is asymptotic[9–11,31]. Our method, which does not require these assumptions, builds on the prediction-based ratio method for rejecting local realism[32]. Applying this method to training data (see below), we obtain a real-valued Bell function T with arguments A, B, X and Y that satisfies $T(A,B,X,Y)>0$ with expectation $\mathbb{E}(T)\leq 1$ for any local-realist distribution that satisfies equation (2). From T we determine the maximum value $1+m$ of $\mathbb{E}(T)$ over all distributions that satisfy equations (2) and (3), where we require that $m>0$. Such a function T induces a Bell inequality (equation (1)) with $\beta=4$ and $s_{xy}^{ab}=T(a,b,x,y)$. Define $T_i=T(A_i,B_i,X_i,Y_i)$ and $V=\prod_{i=1}^{n}T_i$; if the experimenter observes a value of V larger than 1, this indicates a violation of the Bell inequality and the presence of randomness in the data. The randomness is quantified by the 'entropy production theorem' (see below), which we prove in Supplementary Information section 2. We denote all of the settings of both stations with $\mathbf{XY}=X_1Y_1X_2Y_2\ldots X_nY_n$; other sequences such as **AB** and **ABXY** are similarly interleaved over n trials.

The entropy production theorem is as follows. Suppose T is a Bell function that satisfies the above conditions. Then, in an experiment of n trials that obey equations (2) and (3), the following inequality holds for all $\epsilon_p\in(0,1)$ and v_{thresh} satisfying $1\leq v_{\mathrm{thresh}}\leq[1+(3/2)m]^n\epsilon_p^{-1}$:

$$\mathbb{P}_e(\mathbb{P}_e(\mathbf{AB} \mid \mathbf{XY}) > \delta \text{ AND } V \geq v_{\text{thresh}}) \leq \epsilon_{\text{p}} \quad (4)$$

where $\delta = [1 + (1 - \sqrt[n]{\epsilon_{\text{p}} v_{\text{thresh}}})/(2m)]^n$ and $\mathbb{P}_e$ denotes the probability distribution conditioned on the event $E = e$, where e is arbitrary. The expression $\mathbb{P}_e(\mathbf{AB} \mid \mathbf{XY})$ denotes the random variable that takes the value $\mathbb{P}_e(\mathbf{AB} = \mathbf{ab} \mid \mathbf{XY} = \mathbf{xy})$ when $\mathbf{ABXY}$ takes the value $\mathbf{abxy}$.

In words, this theorem says that, with high probability, if V is at least as large as v_{thresh}, then the output $\mathbf{AB}$ is unpredictable, in the sense that no individual outcome $\mathbf{AB} = \mathbf{ab}$ occurs with probability higher than δ, even given the information $\mathbf{XY}E = \mathbf{xy}e$. The theorem supports a protocol that aborts if V takes a value less than v_{thresh}, and passes otherwise. If the probability of passing were 1, then $-\log_2(\delta)$ would be a so-called 'smooth min-entropy'[6]—a quantity that characterizes the number of uniformly distributed bits of randomness that are in principle available in $\mathbf{AB}$. We show in Supplementary Information section 3 that, for constant ϵ_{p}, $-\log_2(\delta)$ is proportional to the number of trials. The number of bits that we can actually extract depends on ϵ_{fin}, the maximum allowed distance of the final output from uniform. We also show in Supplementary Information section 2 that the entropy production theorem can be proved even if the settings probabilities are not known exactly.

To extract the available randomness in $\mathbf{AB}$, we use the TMPS algorithm to obtain an extractor, specifically a function Ext that takes as inputs the string $\mathbf{AB}$ and a 'seed' bit string $\mathbf{S}$ of length d, where $\mathbf{S}$ is uniform and independent of $\mathbf{ABXY}$. Its output is a bit string of length t. $\mathbf{S}$ can be obtained from d additional instances of the random variables X_i, so equation (2) ensures the independence and uniformity conditions on $\mathbf{S}$ that are needed. For the output to be within a distance ϵ_{fin} of uniform independently of $\mathbf{XY}$ and E, the entropy production and extractor parameters must satisfy the constraints given in the 'protocol soundness theorem', which we prove in Supplementary Information section 5. In the statement of the theorem, the measure of distance used is the total variation distance, which is expressed by the left-hand side of equation (6), and 'pass' is the event that V exceeds v_{thresh}.

The protocol soundness theorem is as follows. Let $0 < \epsilon_{\text{ext}}, \kappa < 1$. Suppose that $\mathbb{P}(\text{pass}) \geq \kappa$ and that the protocol parameters satisfy

$$t + 4\log_2(t) \leq -\log_2(\delta) + \log_2(\kappa) + 5\log_2(\epsilon_{\text{ext}}) - 11 \quad (5)$$

Then, the output $\mathbf{U} = \text{Ext}(\mathbf{AB}, \mathbf{S})$ of the function obtained by the TMPS algorithm satisfies

$$\begin{aligned} &\frac{1}{2}\sum_{\mathbf{u},\mathbf{xys}e} |\mathbb{P}(\mathbf{U} = \mathbf{u}, \mathbf{XYS}E = \mathbf{xys}e \mid \text{pass}) \\ &\quad - \mathbb{P}^{\text{unif}}(\mathbf{U} = \mathbf{u})\mathbb{P}(\mathbf{XY}E = \mathbf{xy}e \mid \text{pass})\mathbb{P}^{\text{unif}}(\mathbf{S} = \mathbf{s})| \\ &\quad \leq \frac{\epsilon_{\text{p}}}{\mathbb{P}(\text{pass})} + \epsilon_{\text{ext}} \end{aligned} \quad (6)$$

where $\mathbb{P}^{\text{unif}}$ denotes the uniform probability distribution.

The number of seed bits d that are required satisfies $d = \mathcal{O}[\log(t)\log(nt/\epsilon_{\text{ext}})^2]$; we provide an explicit bound in Supplementary Information section 4. The protocol soundness theorem enables us to quantify the uniformity of the randomness that is produced with an overall final error parameter of $\epsilon_{\text{fin}} = \max(\epsilon_{\text{p}}/\kappa + \epsilon_{\text{ext}}, \kappa)$. (This choice of error parameter is conservative; see Supplementary Information section 5.) For any probability of passing greater than ϵ_{fin}, the total variation distance from uniform (conditionally on passing) is at most ϵ_{fin}.

We applied our protocol to five datasets using a set-up based on that described previously[13], with improvements described in Methods. Each dataset was collected in 5–10 min. Before starting the protocol, we set aside the first 5×10^6 trials of each dataset as training data, which we used to choose the parameters that are needed by the protocol. With the training data removed, the number n of trials used by the protocol was between 2.5×10^7 and 5.5×10^7 for each dataset. We used the training data to determine a Bell function T with statistically strong violation of local realism on the training data according to the prediction-based ratio method[32]; see Supplementary Information section 3. The function T obtained for the fifth dataset, which was the longest in duration and produced the most randomness, assigned values between 0.927 and 1.004 to the 16 different experimental outcomes. We computed thresholds v_{thresh} so that a sample of n i.i.d. trials from the distribution inferred from the training data would have a high probability of exceeding v_{thresh}.

For the fifth dataset, a sample of n i.i.d. trials from the distribution inferred from the training data would have a probability of approximately 0.99 of exceeding a threshold of $v_{\text{thresh}} = 1.5 \times 10^{32}$. Exceeding this threshold would allow the extraction of 1,024 bits that are uniformly distributed to within $\epsilon_{\text{fin}} = 10^{-12}$, using $\epsilon_{\text{p}} = \kappa^2 = 9.025 \times 10^{-25}$ and $\epsilon_{\text{ext}} = 5 \times 10^{-14}$. These values were chosen on the basis of a numerical study of the constraints on the number t of bits extracted for fixed values of $\epsilon_{\text{fin}} = 10^{-12}$. Running the protocol on the remaining 55,110,210 trials with these parameters, the product $\prod_{i=1}^{n} T_i$ exceeded v_{thresh}, and so the protocol passed. Applying the extractor to the resulting output string $\mathbf{AB}$ with a seed of length $d = 315{,}844$, we extracted 1,024 bits, certified to be uniform to within 10^{-12}, the first ten of which are 1110001001. In Fig. 2 we display the extractable bits for alternative choices of ϵ_{fin} for all five datasets.

For the dataset that produced 1,024 new near random bits, our protocol used 1.10×10^8 uniform bits to choose the settings and 3.16×10^5 uniform bits to choose the seed. The strong extractor property[28] of the TMPS algorithm ensures that the seed bits are still uniform, conditional on passing, so they can be recovered at the end of the protocol for use elsewhere. This is not the case for the settings-choice bits because the probability of passing is less than 1. To reduce the entropy used for the settings, our protocol can be modified to use highly biased settings choices[5]. Reducing settings entropy is not a priority if the settings and seed bits come from a public source of randomness, in which case the output bits can still be certified to be unknown to external observers such as Eve and the current protocol is an effective method for private randomness generation[2,10].

For future work, we hope to take advantage of the adaptive capabilities of the entropy production theorem (Supplementary Information section 2) to compensate for experimental drift dynamically during

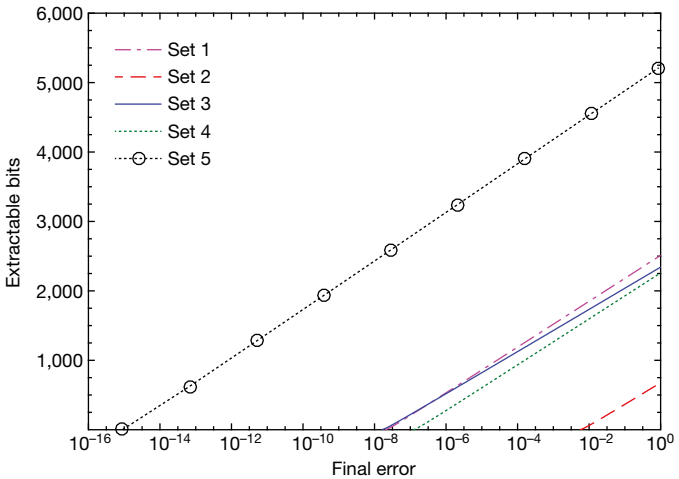

Fig. 2 | Extractable bits as a function of error. The figure shows the trade-off between the final error ϵ_{fin} and the number of extractable bits t for values of v_{thresh} pre-chosen to yield estimated passing probabilities that exceed 95%. These thresholds were met in each case. For all datasets (1–5) we set $\epsilon_{\text{p}} = \kappa^2 = (0.95\epsilon_{\text{fin}})^2$ and $\epsilon_{\text{ext}} = 0.05\epsilon_{\text{fin}}$, a split that was generally found to be near-optimal when numerically maximizing t in equation (5) for fixed values of ϵ_{fin}. The number of trials for datasets 1–5 were $n_1 = 24{,}865{,}320$, $n_2 = 24{,}809{,}970$, $n_3 = 24{,}818{,}959$, $n_4 = 24{,}846{,}822$ and $n_5 = 55{,}110{,}210$.

run time. In view of advances towards practical quantum computing, it is desirable to study the protocol when experimental devices may have long-term quantum memories and remain entangled with Eve after the protocol has begun. This may require more conservative randomness generation.

With the advent of loophole-free Bell tests, we have demonstrated that it is possible to build quantum devices that exploit quantum non-locality to remove many of the device-dependent assumptions in current technological implementations of random-number generators. Generators such as ours provide the best method currently known for physically producing randomness, thereby improving the security of a wide range of applications.

Online content

Any Methods, including any statements of data availability and Nature Research reporting summaries, along with any additional references and Source Data files, are available in the online version of the paper at https://doi.org/10.1038/s41586-018-0019-0.

Received: 6 April 2017; Accepted: 14 February 2018;
Published online 11 April 2018.

1. Acín, A. & Masanes, L. Certified randomness in quantum physics. *Nature* **540**, 213–219 (2016).
2. Pironio, S. & Massar, S. Security of practical private randomness generation. *Phys. Rev. A* **87**, 012336 (2013).
3. Miller, C. A. & Shi, Y. Robust protocols for securely expanding randomness and distributing keys using untrusted quantum devices. *J. ACM* **63**, 33 (2016).
4. Colbeck, R. & Kent, A. Private randomness expansion with untrusted devices. *J. Phys. A* **44**, 095305 (2011).
5. Pironio, S. et al. Random numbers certified by Bell's theorem. *Nature* **464**, 1021–1024 (2010).
6. Vazirani, U. & Vidick, T. Certifiable quantum dice - or, exponential randomness expansion. In *STOC'12 Proc. 44th Annual ACM Symposium on Theory of Computing* (ed. Pitassi, T.) 61–76 (2012).
7. Fehr, S., Gelles, R. & Schaffner, C. Security and composability of randomness expansion from Bell inequalities. *Phys. Rev. A* **87**, 012335 (2013).
8. Chung, K.-M., Shi, Y. & Wu, X. Physical randomness extractors: generating random numbers with minimal assumptions. Preprint at https://arxiv.org/abs/1402.4797 (2014).
9. Nieto-Silleras, O., Pironio, S. & Silman, J. Using complete measurement statistics for optimal device-independent randomness evaluation. *New J. Phys.* **16**, 013035 (2014).
10. Bancal, J.-D., Sheridan, L. & Scarani, V. More randomness from the same data. *New J. Phys.* **16**, 033011 (2014).
11. Thinh, L., de la Torre, G., Bancal, J.-D., Pironio, P. & Scarani, V. Randomness in post-selected events. *New J. Phys.* **18**, 035007 (2016).
12. Hensen, B. et al. Loophole-free Bell inequality violation using electron spins separated by 1.3 kilometres. *Nature* **526**, 682–686 (2015).
13. Shalm, L. K. et al. Strong loophole-free test of local realism. *Phys. Rev. Lett.* **115**, 250402 (2015).
14. Giustina, M. et al. Significant-loophole-free test of Bell's theorem with entangled photons. *Phys. Rev. Lett.* **115**, 250401 (2015).
15. Paar, C. & Pelzl, J. *Understanding Cryptography* (Springer, Heidelberg, 2010).
16. Fischer, M. J., Iorga, M. & Peralta, R. A public randomness service. In *Proc. International Conference on Security and Cryptography (SECRYPT 2011)* (eds Lopez, J. & Samarati, P.) 434–438 (2011).
17. Bell, J. S. On the Einstein Podolsky Rosen paradox. *Phys. Phys. Fiz.* **1**, 195–200 (1964).
18. Bell, J. S., Shimony, A., Horne, M. A. & Clauser, J. F. An exchange on local beables. *Dialectica* **39**, 85–96 (1985).
19. Mayers, D. & Yao, A. Quantum cryptography with imperfect apparatus. In *FOCS'98 Proc. 39th Annual Symposium on Foundations of Computer Science* (ed Motwani, R.) 503–509 (1998).
20. Acín, A. et al. Device-independent security of quantum cryptography against collective attacks. *Phys. Rev. Lett.* **98**, 230501 (2007).
21. Liu, Y. et al. High-speed device-independent quantum random number generation without a detection loophole. *Phys. Rev. Lett.* **120**, 010503 (2018).
22. Rosenfeld, W. et al. Event-ready Bell test using entangled atoms simultaneously closing detection and locality loopholes. *Phys. Rev. Lett.* **119**, 010402 (2017).
23. Abellán, C. et al. Challenging local realism with human choices. *Nature* (in the press).
24. Brunner, N., Cavalcanti, D., Pironio, S., Scarani, V. & Wehner, S. Bell nonlocality. *Rev. Mod. Phys.* **86**, 419–478 (2014).
25. Cirel'son, B. S. Quantum generalizations of Bell's inequality. *Lett. Math. Phys.* **4**, 93–100 (1980).
26. Navascués, M., Pironio, S. & Acín, A. A convergent hierarchy of semidefinite programs characterizing the set of quantum correlations. *New J. Phys.* **10**, 073013 (2008).
27. Trevisan, L. Extractors and pseudorandom generators. *J. Assoc. Comput. Mach.* **48**, 860–879 (2001).
28. Mauerer, W., Portmann, C. & Scholz, V. B. A modular framework for randomness extraction based on Trevisan's construction. Preprint at https://arxiv.org/abs/1212.0520 (2012).
29. Coudron, M. & Yuen, H. Infinite randomness expansion with a constant number of devices. In *STOC'14 Proc. 46th Annual ACM Symposium on Theory of Computing* (ed Shmoys, D.) 427–436 (2014).
30. Arnon-Friedman, R., Dupuis, F., Fawzi, O., Renner, R. & Vidick, T. Practical device-independent quantum cryptography via entropy accumulation. *Nat. Commun.* **9**, 459 (2018).
31. Miller, C. & Shi, Y. Universal security for randomness expansion from the spot-checking protocol. *SIAM J. Comput.* **46**, 1304–1335 (2017).
32. Zhang, Y., Glancy, S. & Knill, E. Asymptotically optimal data analysis for rejecting local realism. *Phys. Rev. A* **84**, 062118 (2011).

Acknowledgements We thank C. Miller and K. Coakley for comments on the manuscript. A.M. acknowledges financial support through NIST grant 70NANB16H207. This work is a contribution of the National Institute of Standards and Technology and is not subject to US copyright.

Reviewer information *Nature* thanks S. Pironio and the other anonymous reviewer(s) for their contribution to the peer review of this work.

Author contributions P.B. led the project and implemented the protocol. P.B., E.K., S.G. and Y.Z. developed the protocol theory. A.M., S.J., A.R. and Y.-K.L. were responsible for the extractor theory and implementation. B.C., S.W.N., M.J.S. and L.K.S. collected and interpreted the data. P.B., E.K., S.G. and L.K.S. wrote the manuscript.

Competing interests The authors declare no competing interests.

Additional information
Supplementary information is available for this paper at https://doi.org/10.1038/s41586-018-0019-0.
Reprints and permissions information is available at http://www.nature.com/reprints.
Correspondence and requests for materials should be addressed to P.B.

METHODS

We used polarization-entangled photons generated by a nonlinear crystal pumped by a pulsed, picosecond laser at approximately 775 nm in a configuration similar to that reported previously[13], but with several improvements to increase the rate of randomness extraction. The repetition rate of the laser was 79.3 MHz and each pulse that entered the crystal had a probability of approximately 0.003 of creating an entangled photon pair in the state $|\psi\rangle \approx 0.982|HH\rangle + 0.191|VV\rangle$ at a centre wavelength of 1,550 nm. By pumping the crystal with approximately five times as much power, and using a 20-mm-long crystal, we were able to increase the per-pulse probability of generating a down-conversion event substantially compared with the previous configuration[13] while maintaining similar overall system efficiencies. The two entangled photons from each pair were sent separately to one of the two measurement stations, which were 187 ± 1 m apart. At Alice and Bob, a Pockels cell and a polarizer combined to allow the rapid switching of measurement bases and the measurement of the polarization state of the incoming photons. Alice's computed optimal polarization measurement angles, relative to a vertical polarizer, were $a = -3.7°$ and $a' = 23.6°$, and Bob's were $b = 3.7°$ and $b' = -23.6°$. Each Pockels cell operated at a rate of 100 kHz, allowing us to perform 100,000 trials per second (the driver electronics on the Pockels cells sets this rate). A 10-MHz oscillator kept Alice's and Bob's time-tagger clocks locked. After passing through the polarization optics, the photons were each coupled into a single-mode fibre and detected using superconducting single-photon nanowire detectors, with Bob's detector operating at approximately 90% efficiency and Alice's detector operating with approximately 92% efficiency[33]. For this experiment, the total symmetric system heralding efficiency was $75.5\% \pm 0.5\%$, which is greater than the 71.5% threshold that is required to close the detection loophole for our experimental configuration after accounting for unwanted background counts at our detectors and slight imperfections in our state-preparation and measurement components.

With this configuration, Bob completed his measurement 294.4 ± 3.7 ns before a hypothetical switching signal travelling at light speed from Alice's Pockels cell could arrive at his station. Similarly, Alice completed her measurement 424.2 ± 3.7 ns before such a signal from Bob's Pockels cell could arrive at her location. The outcome values for each trial were obtained by aggregating the photon detection or non-detection events from several short time intervals, each lasting 1,024 ps and timed to correspond to one pulse of the pump laser. If any photons were detected in the short intervals, then the outcome was '+'; if no photons were detected, then the outcome was '0'. The previous experiment[13] used at most 7 short intervals, but here we were able to include 14 intervals while maintaining space-like separation, which further increased the probability of observing a photon during each trial. For demonstration purposes, Alice and Bob each used Python's random.py module with the default generator (the Mersenne twister) to pick their settings at each trial. This pseudorandom source is predictable, and for secure applications of the protocol in an adversarial scenario, such as if the photon pair source or measurement devices are obtained from an untrusted provider, settings choices must be based on random sources that are effectively not predictable. However, from our knowledge of device construction, we know that our devices have no physical resources for predicting pseudorandom numbers and expect that the measurement settings were effectively independent of the relevant devices so that equations (2) and (3) still hold. We remark that the settings choices for the previous datasets[13] were based on physical random sources.

With the improved detection efficiency, the higher per-trial probability for Alice and Bob to detect a photon, and a higher signal-to-background counts ratio, we are able to improve the magnitude of our Bell violation and to reduce the number of trials that are required to achieve a statistically significant violation by an order of magnitude.

Sample size. No statistical methods were used to predetermine sample size.

Data availability. The photon detection data that support the findings of this study are available in the NIST Published Data Repository (https://doi.org/10.18434/T4/1423448).

33. Marsili, F. et al. Detecting single infrared photons with 93% system efficiency. *Nat. Photon.* **7**, 210–214 (2013).