PHYSICS

An experimental quantum Bernoulli factory

Raj B. Patel^{1*}, Terry Rudolph², Geoff J. Pryde^{1*}

There has been a concerted effort to identify problems computable with quantum technology, which are intractable with classical technology or require far fewer resources to compute. Recently, randomness processing in a Bernoulli factory has been identified as one such task. Here, we report two quantum photonic implementations of a Bernoulli factory, one using quantum coherence and single-qubit measurements and the other one using quantum coherence and entangling measurements of two qubits. We show that the former consumes three orders of magnitude fewer resources than the best-known classical method, while entanglement offers a further fivefold reduction. These concepts may provide a means for quantum-enhanced performance in the simulation of stochastic processes and sampling tasks.

INTRODUCTION

The pursuit of quantum computers (1, 2) has uncovered a number of scenarios where quantum information processing offers a clear advantage over classical means. There exist certain tasks that are intractable using a classical computer but are made possible with quantum computing, supporting the notion of “quantum supremacy” (3). While there are examples where a quantum advantage may exist, unequivocal experimental proof is often unattainable.

A task of interest is that of randomness processing. An example is where a sequence of random variables, obeying a probability distribution, is transformed to produce a new random variable obeying a different probability distribution. Recently, this task has been identified by Dale *et al.* (4) as a basic primitive for which quantum information processing offers advantages over classical stochastic techniques. First, it was shown that the encoding of information on coherent quantum states gives access to a broader range of probability distributions that can be transformed. Second, their approach was shown to require fewer resources, where the resources are akin to the runtime of the task. This randomness processing task has widespread applicability across science and is rooted in processes that are typically simulated by Markov chain Monte Carlo methods. In addition, investigations in this area bear upon our fundamental understanding of quantum randomness (5). In particular, they offer a new avenue for understanding the difference between epistemological classical randomness, owing to noncontextual ignorance about the real state of a system, and quantum randomness, for which no such interpretation is possible.

Here, we present quantum photonic experiments where polarization qubits are used to encode sequences of random variables, which are governed by a probability distribution. From these random variables, algorithmic processing allows for construction of a new distribution while exhibiting a quantum advantage. We show that quantum coherence reduces the resource consumption, or runtime, by several orders of magnitude compared to the best-known classical method, while entanglement offers even further improvements. Before describing the details of our work, we set the scene with some simple examples illustrating the type of processing our work tackles.

Let us first consider the scenario where a fair coin $f(p) = \frac{1}{2}$ is to be simulated from a series of tosses of a classical coin with unknown bias p . Here, p is the probability of a heads outcome, and $p \in (0, 1)$.

Von Neumann’s solution (6) is to toss the coin twice, and if the outcomes are different output the value of the second coin toss, or repeat the procedure if they are the same. Now, suppose the task is to simulate the function $f(p) = p^2$ for $p \in [0, 1]$. This can be achieved by tossing the coin twice. A head is outputted if each individual toss results in heads; otherwise, the output is tails. Some polynomials are well suited to this type of construction. While it is obvious that the function $f(p) = 2p(1 - p)$ may be simulated by tossing a coin twice, the function $f(p) = 3p(1 - p)$ requires noting that $3p(1 - p) = 3p^2(1 - p) + 3p(1 - p)^2$ (7), and as such, the coin must be tossed three times.

In these examples, we have described the scenario of the so-called “Bernoulli factory” (7–16), illustrated in Fig. 1A. Here, one draws from

Fig. 1. The classical Bernoulli factory. (A) Concept of sampling task. A sequence of iid coins, with an unknown bias p , are sampled and processed producing a new coin of bias $f(p)$. (B) Doubling function $f_p(p) = 2p$, as per Eq. 1. The dashed blue plot shows the ideal function, which cannot be constructed classically. The workaround is to truncate the function by ϵ , shown by the solid red line.

¹Centre for Quantum Computation and Communication Technology and Centre for Quantum Dynamics, Griffith University, Brisbane 4111, Australia. ²Department of Physics, Imperial College London, Prince Consort Road, London SW7 2AZ, UK.

*Corresponding author. Email: r.patel@griffith.edu.au (R.B.P.); g.pryde@griffith.edu.au (G.J.P.)

a sequence of independent and identically distributed (iid) Bernoulli random variables (coin tosses), i.e., $\mathbb{P}(X = 0 \equiv \text{Heads}) = p$ and $\mathbb{P}(X = 1 \equiv \text{Tails}) = 1 - p$, for an unknown p . They then process the samples and output a new Bernoulli variable with the probability of obtaining heads $f(p) : (S \subseteq [0, 1]) \rightarrow [0, 1]$. The goal is to not only transform from one probability distribution to another but also to do so with as few coin flips as possible. In the Bernoulli factory, the number of coin flips can be interpreted as the number of times a Bernoulli random variable is queried from a generator, which is proportional to the runtime of the task (13, 14).

These ideas were developed by Keane and O'Brien (8) who derived the necessary and sufficient conditions under which a Bernoulli factory exists for $f(p)$. These conditions are (i) $f(p)$ must be continuous, (ii) $f(p)$ must not approach a value of 0 or 1 exponentially quickly near $p = 0$ or 1, and (iii) $f(p) \neq 0$ or 1 for $p \in (0, 1)$.

RESULTS

Quantum Bernoulli factory for $f(p) = 2p$

Replacing the classical coin with a quantum coin or “quoin,” of the form $|p\rangle = \sqrt{p}|0\rangle + \sqrt{1-p}|1\rangle$ can yield some remarkable advantages (4). The extension to quoins enables algorithmic processing of coherent superpositions and entangled states, with a classical output. We will refer to this as the quantum Bernoulli factory (QBF) and the classical version as the CBF. One interesting feature of the QBF is that an advantage can be gained with quantum coherence alone. The conditions imposed by Keane and O'Brien are now relaxed in the quantum setting, allowing a larger class of functions to be constructed.

The function we choose to study, and perhaps the most important, is the “Bernoulli doubling” function

$$f_{\wedge}(p) = 2p \equiv \begin{cases} 2p & p \in [0, 1/2] \\ 2(1-p) & p \in (1/2, 1] \end{cases} \quad (1)$$

since it serves as a primitive for other factories (9). That is, the ability to sample from this function allows any other analytical function to be constructed that is bounded at less than unity in $(0, 1)$. Notice that this function cannot be constructed classically since $f_{\wedge}(0.5) = 1$ violates condition (iii). In the classical setting, the workaround is to truncate the function by ϵ , where $0 < \epsilon < 1$, such that $f_{\wedge}(p) = 2p \approx \min(2p, 1 - \epsilon)$ (8, 9, 11–13), as shown in Fig. 1B. From (4), the QBF for Eq. 1 can be realized by first rewriting the function as $f_{\wedge}(p) = 1 - \sqrt{1 - 4p(1-p)}$ and performing a series expansion

$$\begin{aligned} f_{\wedge}(p) &= \sum_{k=1}^{k_{\max}} \binom{2k}{k} \frac{1}{(2k-1)2^{2k}} (4p(1-p))^k \\ &= \sum_{k=1}^{k_{\max}} q_k g_k \end{aligned} \quad (2)$$

Here, $k > 0$, q_k is independent of p , and $g_k = (4p(1-p))^k$. Typically, $k_{\max} = \infty$; however, in realistic experimental scenarios, finite $k_{\max}$ values are considered. This representation allows us to reduce the problem to finding a construction for g_k —or k consecutive heads outcomes of tossing a g_1 -coin—where a g_1 -coin is defined as a coin with a bias $g_1(p) = 4p(1-p)$. The main task is thus to efficiently produce such a g_1 -coin. Performing a joint two-qubit measurement on two

p -quoins $|p\rangle \otimes |p\rangle$ in the Bell basis, $\{|\psi^{\pm}\rangle = (|01\rangle \pm |10\rangle)/\sqrt{2}, |\phi^{\pm}\rangle = (|00\rangle \pm |11\rangle)/\sqrt{2}\}$, we find that

$$\begin{aligned} \mathbb{P}(\phi^{-} | (\psi^{+} \cup \phi^{-})) &= (2p-1)^2, \\ \mathbb{P}(\psi^{+} | (\psi^{+} \cup \phi^{-})) &= 4p(1-p) = g_1(p), \\ \mathbb{P}(\phi^{-} | (\psi^{+} \cup \phi^{-})) + \mathbb{P}(\psi^{+} | (\psi^{+} \cup \phi^{-})) &= 1 \end{aligned} \quad (3)$$

where $\mathbb{P}(\psi^{+} | (\psi^{+} \cup \phi^{-}))$ is read as the probability of measuring $|\psi^{+}\rangle$ given that we restrict ourselves to measuring states $|\psi^{+}\rangle$ or $|\phi^{-}\rangle$. The algorithm runs by first generating an index k with a probability q_k . A joint measurement, in a restricted Bell basis, is then performed on two p -quoins. If k consecutive $|\psi^{+}\rangle$ outcomes are obtained, then the toss of an $f_{\wedge}(p)$ -coin is heads; otherwise, if the outcome of the measurement is $|\phi^{-}\rangle$, then the output is tails.

Two-qubit experimental QBF

The required measurements are well suited to our linear optics implementation shown in Fig. 2. A 404 nm, vertically (V) polarized, continuous-wave laser beam pumps a nonlinear BiBO crystal generating a degenerate pair of horizontally (H) polarized photons, $|H\rangle_1 \otimes |H\rangle_2$. The photons are spectrally filtered using longpass filters and 3 nm bandpass filters centered at 808 nm and are sent via single-mode fiber to a Bell-state analyzer. This particular arrangement contains additional motorized half-wave plates (MHWPs), which set the bias value of each p -quoin. It is well known that the standard linear optical Bell-state analyzer (17, 18), relying on Hong-Ou-Mandel interference, is capable of unambiguously discriminating between the $|\psi^{+}\rangle$ and $|\psi^{-}\rangle$ Bell states. We implement an $X_{\frac{\pi}{2}} \otimes X_{-\frac{\pi}{2}}$ operation on the qubits before the measurement using quarter-wave plates (QWPs) at optic axes (OA) $\pm 45^\circ$, which allows the desired states, $|\psi^{+}\rangle$ and $|\phi^{-}\rangle$, to be identified. The photons interfere on a 50:50 nonpolarizing beamsplitter (NPBS), while polarizing beamsplitters (PBSs) enable H- and V-polarized photons to be separated spatially before being detected using single-photon avalanche diodes (APDs). The sequence of detection events is time-tagged, which allows us to identify the exact order in which a $g_1(p)$ -coin toss resulted in a heads ($|\psi^{+}\rangle$) or a tails ($|\phi^{-}\rangle$). The data are postprocessed (see Materials and Methods for details) using a computer, and $f_{\wedge}(p)$ is constructed. Figure 3 (A to D) shows the experimental data (circles) taken for $k_{\max} \in \{1, 10, 100, 2000\}$. We see that the data agree strongly with the ideal theoretical plots (dotted lines). The red curves in each plot represent the expected data based on a model, which takes into account the nonideal splitting ratio of our NPBS, extinction ratios of our polarization optics, and any mode-mismatch in our interferometer. The experimentally measured Hong-Ou-Mandel two-photon interference visibility was found to be $(99.7^{+0.3}_{-1.0})\%$. The experimental data show an excellent agreement with our model. For lower values of k , the data show a more rounded peak near $p = 0.5$, which becomes sharper for larger k . In our experimental run, we were able to generate up to a single $g_{2036}(p)$ coin, i.e., up to 2036 consecutive heads outcomes of the $g_1(p)$ coin. Higher-order $g_k(p)$ coins are more susceptible to small experimental imperfections, which may lead to erroneous coincident detections. For more reliable statistics, and for comparison later on, in Fig. 3D, we restrict the expansion to $k_{\max} = 2000$, where we obtain $f_{\wedge}(0.5) = 0.935 \pm 0.003$. In Fig. 3E, we calculate the mean p -quoin consumption for each $f_{\wedge}(p)$ -coin. Note that increasingly more quoins are required near $p = 0.5$, as we expect. We require an average (over p) of ≈ 11 quoins to construct $f_{\wedge}(p) = 2p$ when using the quantum coherence and entangling measurements of two p -quoins.

Fig. 2. Experimental arrangement for the QBF using joint measurements of two p -quoins. A pair of H-polarized photons are generated via type I down-conversion in a nonlinear BiBO crystal. They are sent (indicated by red arrows) to a Bell-state analyzer arrangement containing additional MHPW, which set the bias value of each p -quoin, and QWPs, one at OA $+45^\circ$ and one at OA -45° , which enables $|\psi^+\rangle$ and $|\phi^-\rangle$ to be identified. The photons interfere on a 50:50 NPBS, while the PBS enable H- and V-polarized photons to be separated spatially before being detected using single-photon APDs. Detection events are time-tagged and analyzed using a computer.

Fig. 3. Experimental data for the two-qubit QBF. The function $f_\Lambda(p) = 2p$ is constructed using joint measurements of two p -quoins for (A) $k_{\max} = 1$, (B) $k_{\max} = 10$, (C) $k_{\max} = 100$, and (D) $k_{\max} = 2000$. The dotted blue lines are the ideal theoretical functions, and the red lines represent a model taking experimental imperfections into consideration. Error bars were too small to be included (see Materials and Methods). (E) Mean p -quoin consumption for $k_{\max} = 2000$.

Single-qubit experimental QBF

We now show how $f_\Lambda(p)$ can be constructed using single-qubit measurements, where we exploit quantum coherence alone. To do so, we use the best-known algorithm for constructing $g_1(p)$ with single-qubit measurements, which was recently demonstrated using superconducting qubits (19). The algorithm makes use of additional, intermediate quoins denoted by q , m , n , s , and t , each with a unique probability distribution, which are derived from p -quoins. Figure 4A illustrates the procedure where red (blue) arrows indicate a heads (tails) outcome of a particular toss. A more thorough description is provided in the Supplementary Materials. To begin with, two p -quoins are produced, the second of which is measured in the X-basis to produce a q -quoin (lower branch). In the upper branch, the p -quoin is tossed twice, and if the outcome is different each time an m -quoin is produced with the outcome heads, otherwise, tails is outputted. Similarly, in the lower branch where a q -quoin is tossed twice with different outcomes, an n -quoin with a heads outcome results. The m - and n -quoins are both tossed twice. In each case, if the first toss results in tails, a new quoin is produced, s or t , with a tails outcome. If, however, the first toss gives heads and the second gives tails, then heads is outputted in each case. Otherwise, the protocol is repeated from the beginning, and two p -quoins are sampled again. Given the successful construction of an s - and t -quoin, if they have the value heads (tails) and tails (heads), respectively, the outcome of g -coin toss is heads (tails). If the outcome is the same each time, the protocol is repeated. From the successful sampling of the g_1 -coin, $f_\Lambda(p)$ can be constructed, as outlined earlier.

The experimental configuration is shown in Fig. 4B. Using the same photon-pair source as before, one photon is used as a herald,

Fig. 4. Construction of $f_\lambda(p) = 2p$ using single-qubit measurements of p -quoins. (A) The algorithm we use (see the Supplementary Materials). The upper (lower) branch begins with the measurement of the p -quoin in the Z-basis (X-basis). Dashed red (blue) arrows indicate a heads (tails) outcome of the quoin toss. Failure to achieve the appropriate outcome requires the protocol to be repeated until success. (B) Experimental arrangement for the QBF using single-qubit measurements of p -quoins. Red arrows indicate photon inputs from the source (not shown). A single photon encounters a MHWP, which sets p . A HWP set to OA enables Z-basis measurements to be performed for each p . The partner photon, which serves as a herald, is detected directly by an APD. Setting the HWP to OA + 22.5° results in an X-basis measurement. Two sets of time-tag data are recorded, allowing p and q -quoins to be sampled.

while the other is sent to an arrangement of MHWP, HWP, and a PBS. Again, the MHWP sets the bias p , while the HWP set to OA (OA + 22.5°) enables Z-basis (X-basis) measurements to be performed for each p . Time-tags are recorded for each measurement basis independently, and the construction of $f_\lambda(p)$ then follows by sampling from the two datasets. Figure 5 (A to D) shows experimental data (circles) taken for $k_{\max} \in \{1, 10, 100, 2000\}$. The data show excellent agreement with theory under ideal conditions, albeit with a slight skew in the data which we attribute to mechanical drift in the fiber coupling. As one might expect, single-qubit measurements, which do not rely on nonclassical interference or multi-qubit coherence, can be performed with higher fidelity than joint measurements on two qubits. As such, in the case of $k_{\max} = 2000$, we obtain $f_\lambda(0.5) = 0.977 \pm 0.006$.

Of particular interest is a comparison of resource consumption between the two QBFs we have presented. For a fair comparison with

the two-qubit QBF, we choose to restrict the series expansion of the single-qubit QBF to $k = 82$, which results in $f_\lambda(0.5) = 0.935 \pm 0.006$. Figure 5E shows the mean p -quoin consumption for each $f_\lambda(p)$ -coin. Averaging overall p , we require ≈ 52 quoins to construct $f_\lambda(p) = 2p$ when using the quantum coherence and single-qubit measurements of p -quoins, which is approximately a fivefold increase in resources over the two-qubit case.

Quantum advantage

Owing to small experimental imperfections, we are unable to exactly achieve $f_\lambda(0.5) = 1$ in our construction of $f_\lambda(p) = 2p$; however, this does provide an avenue for comparing the QBF to the CBF. We can frame the situation as a refereed game played between two parties, the quantum player who has a QBF and a classical player who has a CBF. The referee prepares p -quoins and sends them to the quantum player who is tasked with constructing, or approximating, $f_\lambda(p) = 2p$ as best as they can. The quantum player can request a large, albeit finite, number of quoins. Their results—the quantum-obtained approximation to $f_\lambda(p)$, for a range of p values—are sent to the classical player who must reproduce it using fewer resources. In the game, the quantum player achieves $f_\lambda(0.5) = 1 - \epsilon$. The classical player's strategy is as follows. First, they perform a least-squares fit of the data, for all p , using a positively weighted sum of Bernstein polynomials

$$\tilde{f}(p) = D \sum_{j=0}^N \frac{A_j}{D} \binom{N}{j} p^j (1-p)^{N-j} = D.H(p) \quad (4)$$

where $D = \sum_{j=0}^N A_j$ and $D.H(p) \leq 1 - \epsilon$. The parameters $A_j \geq 0$ are fitting parameters (see the Supplementary Materials) and $\epsilon = 1 - f_\lambda(0.5)$. As with the previously mentioned examples, polynomials of the form $p^j(1-p)^{N-j}$ require N coins for exact sampling (14). This approach takes into consideration the nuances of the experimental data, which deviate from the ideal truncated function shown in Fig. 1B. It then follows from (14) that the mean coin consumption is

$$\bar{N}_c \sim \frac{9.5DN}{\epsilon} \quad (5)$$

To determine the optimal N , the classical player performs an optimization routine where the R-squared value is maximized for a range of N . For the data in Fig. 3D, $\epsilon = 0.065$, $N = 27$, and $D = 14.17$ (see the Supplementary Materials), resulting in $\bar{N}_c \sim 56126$ coins, on average, which is three orders of magnitude greater than the quantum player, who thus wins the game. To the best of our knowledge, this is the optimum strategy that the classical player can use.

Last, we remark on how the resource consumption scales with ϵ . From (14), it was shown that the classical coin consumption for the truncated function shown in Fig. 1B is given by $\bar{N}_c \sim 19\epsilon^{-1}$. Taking into consideration the two-qubit QBF, we calculate the mean p -quoin consumption for a range of ϵ . The two-qubit QBF presented here shows an improvement where the mean quoin consumption scales as $\bar{N}_q \sim 2\epsilon^{-0.5}$, which is in broad agreement with the scaling derived from the experimental data, $\bar{N}_q \sim 3\epsilon^{-0.4}$. As expected, this further supports the notion of a quantum advantage in resource consumption over the best-known classical algorithm.

DISCUSSION

The Bernoulli factory offers a fresh perspective from which information processing can be enhanced by quantum physics. Specifically, we have

Fig. 5. Experimental data for the single-qubit QBF. The function $f_k(p) = 2p$ is constructed using single-qubit measurements of p -quoins for (A) $k_{\max} = 1$, (B) $k_{\max} = 10$, (C) $k_{\max} = 100$, and (D) $k_{\max} = 2000$. The dotted blue lines are the ideal theoretical functions. Error bars were too small to be included (see Materials and Methods). (E) Mean p -coin consumption for $k_{\max} = 82$.

experimentally demonstrated a quantum advantage in the processing of randomness in a QBF under two different scenarios. Our work confirms that quantum coherence can provide a large reduction (three orders of magnitude) in resources over the CBF and that quantum entanglement provides a further fivefold reduction. An interesting but large and challenging theoretical extension would be to compare the resources for a CBF and a QBF to construct an approximation, $j(p)$, to a general function $f(p)$, to a certain degree of accuracy given some suitable measure between functions.

In addition, while our implementation uses bipartite entanglement, an interesting question is: How does this advantage scale when considering multipartite entangled systems? The QBF described here takes iid quoins as its input and outputs a coin. Lifting these restrictions, allowing quoins to be outputted rather than just coins is expected to give rise to other classes of factories and constructible functions (20, 21).

In the study of quantum modeling protocols, the encoding of the input information with quantum states demonstrating coherence has also shown the potential to yield resource advantages in terms of memory (22–25). In particular, the QBF has recently drawn comparisons to the quantum transducer (23), which is a model of an input-output process requiring a lesser amount of past knowledge and complexity compared to its classical counterpart to simulate the future state of the system. Further investigation is required to determine whether the QBF can offer additional insight into the study of processes that have a causal dependence.

MATERIALS AND METHODS

Data processing

For each setting of p , $\sim 3 \times 10^6$ time tags were recorded. Since two photons are always created at the same time during the down-conversion process, the data were filtered, such that only tags

occurring within a window of 6.25 ns remained. This process eliminates most of the spurious events due to ambient light or detector dark counts. The data were traversed, and a tally was kept of the number of coincidence detections corresponding to a heads outcome for each of the $g_k(p) \equiv g_1^k(p)$ coins as well as the number of tails outcomes. $g_k(p)$ was then calculated as $\#heads_k / (\#heads_k + \#tails_k)$. We also note that a tally was kept of the total number of p -quoins required to produce an $f_k(p) = 2p$ coin, which also includes cases where coincident detections correspond to neither a head nor a tail $g_1(p)$ -coin. These events may occur due to imperfect extinction of the polarization optics or the finite polarization dependence of the nonpolarizing optics. This total is then weighted by q_k , allowing the mean p -coin consumption to be calculated.

Poissonian uncertainties arose because we counted a large number of g_k -coins within a fixed data collection time window. Errors quoted throughout the main text were calculated, assuming Poissonian counting statistics of the coincidence detections, which are integrated to give $\#heads_k$ and $\#tails_k$. Errors in $f_k(p)$ typically varied between $\pm 2 \times 10^{-4}$ and $\pm 3 \times 10^{-3}$.

SUPPLEMENTARY MATERIALS

Supplementary material for this article is available at <http://advances.sciencemag.org/cgi/content/full/5/1/eaau6668/DC1>

Section S1. Constructing $g_1(p)$ in the single-qubit QBF

Section S2. Bernstein polynomial fit of the data

Fig. S1. Least-squares fit of $f_k(p) = 2p$.

REFERENCES AND NOTES

1. T. D. Ladd, F. Jelezko, R. Laamme, Y. Nakamura, C. Monroe, J. L. O'Brien, Quantum computers. *Nature* **464**, 45–53 (2010).
2. R. B. Patel, J. Ho, F. Ferreyrol, T. C. Ralph, G. J. Pryde, A quantum Fredkin gate. *Sci. Adv.* **2**, e1501531 (2016).
3. A. W. Harrow, A. Montanaro, Quantum computational supremacy. *Nature* **549**, 203–209 (2017).
4. H. Dale, D. Jennings, T. Rudolph, Provable quantum advantage in randomness processing. *Nat. Commun.* **6**, 8203 (2015).
5. P. Grangier, A. Auffèves, What is quantum in quantum randomness? arXiv:1804.04807 (2018).
6. J. von Neumann, Various techniques used in connection with random digits. *Appl. Math. Ser.* **12**, 36–38 (1951).
7. J. Wästlund, "Functions arising by coin flipping," (Technical Report, KTH, Stockholm, 1999).
8. M. S. Keane, G. L. O'Brien, A Bernoulli factory. *ACM Trans. Model. Comput. Simul.* **4**, 213–219 (1994).
9. Ş. Nacu, Y. Peres, Fast simulation of new coins from old. *Annal. Appl. Probab.* **15**, 93–115 (2005).
10. E. Mossel, Y. Peres, C. Hillar, New coins from old: COmputing with unkown bias. *Combinatorica* **25**, 707–724 (2005).
11. K. Łatuszyński, I. Kosmidis, O. Papaspiliopoulos, G. O. Roberts, Simulating events of unknown probabilities via reverse time martingales. *Random Struct. Algor.* **38**, 441–452 (2011).
12. A. C. Thomas, J. H. Blanchet, A practical implementation of the Bernoulli Factory. arXiv:1106.2508 (2011).
13. J. M. Flegel, R. Herbei, Exact sampling for intractable probability distributions via a Bernoulli factory. *Electron. J. Stat.* **6**, 10–37 (2012).
14. M. Huber, Nearly optimal Bernoulli factories for linear functions. *Comb. Probab. Comput.* **25**, 577–591 (2016).
15. M. Huber, Optimal linear Bernoulli factories for small mean problems. *Method. Comput. Appl. Probab.* **19**, 631–645 (2017).
16. S. Asmussen, P. W. Glynn, H. Thorisson, Stationarity detection in the initial transient problem. *ACM Trans. Model. Comput. Simul.* **2**, 130–157 (1992).
17. K. Mattle, H. Weinfurter, P. G. Kwiat, A. Zeilinger, Dense coding in experimental quantum communication. *Phys. Rev. Lett.* **76**, 4656–4659 (1996).
18. D. Bouwmeester, J.-W. Pan, K. Mattle, M. Eibl, H. Weinfurter, A. Zeilinger, Experimental quantum teleportation. *Nature* **390**, 575–579 (1997).

19. X. Yuan, K. Liu, Y. Xu, W. Wang, Y. Ma, F. Zhang, Z. Yan, R. Vijay, L. Sun, X. Ma, Experimental quantum randomness processing using superconducting qubits. *Phys. Rev. Lett.* **117**, 010502 (2016).
20. J. Jiang, J. Zhang, X. Sun, Quantum-to-quantum Bernoulli factory problem. *Phys. Rev. A* **97**, 032303 (2018).
21. T. J. Yoder, Building and bounding quantum Bernoulli factories (2015); www.scottaaronson.com/6s899/tedyoder.pdf.
22. M. Gu, K. Wiesner, E. Rieper, V. Vedral, Quantum mechanics can reduce the complexity of classical models. *Nat. Commun.* **3**, 762 (2012).
23. J. Thompson, A. J. P. Garner, V. Vedral, M. Gu, Using quantum theory to simplify input–output processes. *npj Quantum Inf.* **3**, 6 (2017).
24. J. Thompson, A. J. P. Garner, J. R. Mahoney, J. P. Crutchfield, V. Vedral, M. Gu, Causal asymmetry in a quantum world. *Phys. Rev. X* **8**, 031013 (2018).
25. C. Aghamohammadi, S. P. Loomis, J. R. Mahoney, J. P. Crutchfield, Extreme quantum memory advantage for rare-event sampling. *Phys. Rev. X* **8**, 011025 (2018).

Acknowledgments: We acknowledge T. Yoder for useful discussions regarding the resource consumption in the CBF and QBF and S. Wollmann for early contributions

to the project. **Funding:** This work was supported by the Australian Research Council Centre of Excellence for Quantum Computation and Communication Technology (project numbers CE110001027 and CE170100012) and the Engineering and Physical Sciences Research Council. **Author contributions:** R.B.P. and G.J.P. designed the experiments, and R.B.P. performed them. T.R. provided theory support. R.B.P. analyzed the data, with input from T.R. and G.J.P. R.B.P. wrote the manuscript, with input from the other authors. **Competing interests:** The authors declare that they have no competing interests. **Data and materials availability:** All data needed to evaluate the conclusions in the paper are present in the paper and/or the Supplementary Materials. Additional data related to this paper may be requested from the authors.

Submitted 4 July 2018

Accepted 11 December 2018

Published 25 January 2019

10.1126/sciadv.aau6668

Citation: R. B. Patel, T. Rudolph, G. J. Pryde, An experimental quantum Bernoulli factory. *Sci. Adv.* **5**, eaau6668 (2019).