Supplementary Information for "Unconditional and exponentially large violation of classicality"

Marcello Benedetti[1], Gabriel Marin-Sanchez[1], Jordi Weggemans[1], Matthias Rosenkranz[1], and Harry Buhrman[1,2]

[1]Quantinuum, London, United Kingdom
[2]University of Amsterdam & QuSoft, The Netherlands

Contents

Supplementary Note 1: Bell functional and violation of classicality

The complement sampling game is a single-player game that does not rely on entanglement as a resource. Regardless, it is insightful to define the game analogously to non-local games, and equip it with Bell-like inequalities. We define a Bell functional [1] for the game as the expected score over the randomness of input and output, i.e. $\mathbb{E}_{\text{input}}\mathbb{E}_{\text{output}|\text{input}}[\text{score}]$. Importantly, the score function is chosen so that if the player were to play completely at random, the expected score would be zero. This allows us to compare strategies in terms of their advantage relative to random guessing.

General complement sampling functional. Let ω be a non-empty and finite universe, and $\mathcal{P}(\omega)$ be its power-set. The referee draws an instance $S \in \mathcal{P}(\omega) \setminus \{\varnothing, \omega\}$ with probability $p(S)$ and provides the player with the subset state $|S\rangle = \frac{1}{\sqrt{|S|}} \sum_{x \in S} |x\rangle$ as input. The player adopts any strategy they want and outputs a bit string in ω. We define a strategy as a function $h : \Sigma^n \to \Gamma^n$, from the set of quantum states over n qubits, Σ^n, to the set of probability vectors over n bits, Γ^n. We use $h(y \mid S)$ to indicate the probability of output y given input $|S\rangle$. The Bell functional value of this game is the expected score

$$V(p, h) = \sum_{S \in \mathcal{P}(\omega) \setminus \{\varnothing, \omega\}} p(S) \sum_{y \in \omega} h(y \mid S)\, \sigma\left(S, y\right), \tag{1}$$

where the ± 1 score is

$$\sigma(S, y) = \begin{cases} +1, & y \in \bar{S}, \\ -1, & y \notin \bar{S}. \end{cases} \tag{2}$$

Equivalently, $\sigma(S, y) = 2 \left[\!\left[y \in \bar{S} \right]\!\right] - 1$ where $[\![\cdot]\!]$ is the Iverson bracket. For the purpose of showing an exponentially large violation of classicality of our game, we can restrict p to be the uniform distribution over a family of subsets $\mathcal{F} \subset P(\omega)$. We denote the corresponding value of the strategy with $V(\mathcal{F}, h)$. Given a set of strategies H, we define its value as

$$V(\mathcal{F}, H) = \sup_{h \in H} V(\mathcal{F}, h).$$

We are interested in computing the ratio

$$\frac{V(\mathcal{F}, Q) - V(\mathcal{F}, h_{\text{rnd}})}{V(\mathcal{F}, C) - V(\mathcal{F}, h_{\text{rnd}})},$$

for specific families $\mathcal{F}$, quantum strategies Q, and classical strategies C. Here h_{rnd} is the strategy that guesses y completely at random, and since $V(\mathcal{F}, h_{\text{rnd}}) = 0$, we hereafter drop it from the ratio. This is without loss of generality since the ratio above is invariant to affine transformations of the score $\sigma' = a\sigma + b$.

A quantum strategy can manipulate $|S\rangle$ coherently, and perform arbitrary measurements to produce the output distribution. In contrast, a classical strategy begins with a measurement of $|S\rangle$ in the computational basis, then processes the result on a classical computer. In this latter case, the input model is equivalent to directly receiving one bit string sampled uniformly from S.

Special case: Bernstein-Vazirani subset states. Let us specialize the above definitions to Bernstein-Vazirani (BV) subsets:

$$\begin{aligned}
&\omega = \{0,1\}^n && |\omega| = 2^n \\
&\mathcal{F}_{\text{BV}} = \{S_{u,b} \mid u \in \omega \setminus \{0\}^n \text{ and } b \in \{0,1\}\} && |\mathcal{F}_{\text{BV}}| = 2(2^n - 1) \\
&S_{u,b} = \{x \in \omega \mid u \cdot x \equiv b \pmod 2\} && |S_{u,b}| = 2^{n-1}.
\end{aligned}$$

We let the input follow a uniform probability distribution over $\mathcal{F}_{\text{BV}}$

$$p(S) = \begin{cases} \frac{1}{2(2^n-1)}, & S \in \mathcal{F}_{\text{BV}}, \\ 0, & \text{otherwise.} \end{cases} \tag{3}$$

Quantum value. Since any BV subset contains half the elements of the universe, a simple application of the swapper circuit gives the complement subset state $\left|\bar{S}\right\rangle = (2\left|+^n\right\rangle\!\left\langle+^n\right| - I)\left|S\right\rangle$ [2]. A measurement in the computational basis gives a bit string from the complement. This quantum strategy implements the conditional probability distribution

$$h(y \mid S) = \begin{cases} \frac{1}{2^{n-1}}, & y \in \bar{S} \\ 0, & \text{otherwise.} \end{cases} \tag{4}$$

It is easy to verify that in the quantum case plugging in Eqs. (3) and (4) in Eq. (1) gives an expected score of one, which is optimal. Thus for the set of all quantum strategies Q we have $V(\mathcal{F}_{\text{BV}}, Q) = 1$.

Classical value. The classical strategy begins by measuring the subset state in the computational basis, thus reading off a bit string x with Born probability

$$h(x \mid S) = |\left\langle x|S\right\rangle|^2 = \begin{cases} \frac{1}{2^{n-1}}, & x \in S \\ 0, & \text{otherwise.} \end{cases} \tag{5}$$

Then the strategy outputs a bit string y with probability $h(y \mid x)$. For example, consider a function s that outputs a particular bit string $s(x) \neq x$ (i.e. it does not have any fixed point). Such function leads to a deterministic output

$$h(y \mid x) = [\![\, y = s(x) \,]\!]\,. \tag{6}$$

The two steps just described implement

$$\begin{aligned} h(y \mid S) &= \sum_{x \in \omega} h(y \mid x)\, h(x \mid S) \\ &= \frac{1}{2^{n-1}} \sum_{x \in S} [\![\, y = s(x) \,]\!]\,, \end{aligned} \tag{7}$$

where we marginalize over the measurement outcomes and use conditional independence $h(y|x, S) = h(y|x)$ since, under the assumption of a trusted referee, no information about the subset could have leaked. We now analyze this classical strategy, and later we will justify why it is essentially the optimal one. Plugging Eqs. (3) and (7) in Eq. (1) we obtain the expected score

$$\begin{aligned} V(\mathcal{F}_{\mathrm{BV}}, h) &= \sum_{S \in \mathcal{F}_{\mathrm{BV}}} \frac{1}{2(2^n - 1)} \sum_{y \in \omega} \frac{1}{2^{n-1}} \sum_{x \in S} [\![\, y = s(x) \,]\!]\, \sigma\,(S, y) \\ &= \frac{1}{(2^n - 1)2^n} \sum_{S \in \mathcal{F}_{\mathrm{BV}}} \sum_{x \in S} \sigma\,(S, s(x)) \\ &= \frac{1}{(2^n - 1)2^n} \sum_{u \in \omega \setminus \{0\}^n} \sum_{b \in \{0,1\}} \sum_{x \in S_{u,b}} \sigma\,(S_{u,b}, s(x)) \\ &= \frac{1}{(2^n - 1)2^n} \sum_{u \in \omega \setminus \{0\}^n} \sum_{x \in \omega} \sigma\,(S_{u,u\cdot x}, s(x))\,. \end{aligned}$$

In the third line we express the subsets in terms of their parameters u and b. In the fourth line we use that for fixed u, any x is either in $S_{u,0}$ or in $S_{u,1}$ depending on the result of $u \cdot x \pmod 2$. Next we rewrite the score as

$$\sigma\,(S_{u,u\cdot x}, s(x)) = 2\,[\![\, u \cdot s(x) \neq u \cdot x \,]\!] - 1,$$

so that

$$V(\mathcal{F}_{\mathrm{BV}}, h) = \left(\frac{2}{(2^n - 1)2^n} \sum_{x \in \omega} \sum_{u \in \omega \setminus \{0\}^n} [\![\, u \cdot s(x) \neq u \cdot x \,]\!] \right) - 1. \tag{8}$$

We now count how many u's satisfy the condition $u \cdot s(x) \neq u \cdot x$ or, equivalently, $u \cdot (s(x) \oplus x) = 1$. Recall that we assumed a map without fixed points, i.e. $s(x) \neq x$ for all x. Letting $d(x) := s(x) \oplus x$, our assumption gives $d(x) \neq \{0\}^n$ for all x. For fixed x, exactly half of the possible bit strings u satisfy $u \cdot d(x) = 1$. Therefore each x contributes $2^n/2$ to the summation in Eq. (8). We arrive at

$$V(\mathcal{F}_{\mathrm{BV}}, h) = \left(\frac{2}{(2^n - 1)2^n} \sum_{x \in \omega} \frac{2^n}{2} \right) - 1 = \frac{2^n}{2^n - 1} - 1 = \frac{1}{2^n - 1}.$$

We remark that: i) any fixed-point-free map s (e.g. $s(x) = x \oplus e_1$ for a fixed nonzero e_1) attains this value; ii) maps with fixed points $s(x) = x$ can only perform worse since x is always sampled from S following Eq. (5); and iii) randomization cannot improve the value since each x contributes the same amount regardless of which $s(x) \neq x$ is chosen. We conclude that the optimal expected score over classical strategies is

$$V(\mathcal{F}_{\mathrm{BV}}, C) = \max_{h \in C} V(\mathcal{F}_{\mathrm{BV}}, h) = \frac{1}{2^n - 1} \tag{9}$$

and the maximum is attained by the classical strategy described above.

Quantum-classical ratio. Using the above results for BV subset states we arrive at the ratio of the expected scores of quantum and classical strategies

$$\frac{V(\mathcal{F}_{\mathrm{BV}}, Q)}{V(\mathcal{F}_{\mathrm{BV}}, C)} = 2^n - 1.$$

Our complement sampling game gives an unconditional and exponentially large violation of classicality.

Supplementary Note 2: Classical sample complexity bounds for Bernstein-Vazirani subsets

For the subset state family $\mathcal{F}_{\mathrm{BV}}$, we will show in this Supplementary Note that the classical randomized sample complexity to solve complement sampling is $\Theta(n)$.

Upper bound. The key idea of our upper bound procedure is that, for any fixed $x_0 \in S$ and any uniformly random $x_j \in S$, the vector $d_j = x_j \oplus x_0$ is uniformly random in the $(n-1)$-dimensional subspace $V = \{v \in \mathbb{F}_2^n \mid v \cdot u \equiv b \pmod 2\}$. If the matrix D whose rows are the vectors $\{d_j\}$ spans V, then $\ker(D) = \mathrm{span}\{u\}$. Thus u is the unique nonzero null vector and it can be found by Gaussian elimination. Once we have found u, we can use x_0 to check whether $b = 0$ or $b = 1$, and then sample uniformly random bit strings y until we find one that satisfies $y \cdot u \equiv 1 - b \pmod 2$. This last step requires only two attempts in expectation and does not cost new input samples (we sample from $\{0,1\}^n$ and not from S). Returning y solves complement sampling with $\mathcal{O}(n)$ sample complexity.

Theorem 1. *Fix any $S \in \mathcal{F}_{\mathrm{BV}}$. For any $\epsilon > 0$, there is a classical procedure that uses*

$$m = n + \lceil \log(1/\epsilon) \rceil$$

i.i.d. samples from S (with replacement) and recovers the hidden $u \in \{0,1\}^n \setminus \{0\}^n$ with probability at least $1 - \epsilon$. The runtime is polynomial in n and $\log(1/\epsilon)$.

Proof. Let $u \in \{0,1\}^n \setminus \{0\}^n$ and $b \in \{0,1\}$ be such that

$$S = \{x \in \{0,1\}^n \mid x \cdot u \equiv b \pmod 2\}.$$

Repeatedly measuring different copies of $|S\rangle$ in the computational basis is equivalent to sampling $x_0, x_1, \ldots, x_{m-1}$ i.i.d. uniformly from S with replacement. For $j \in \{1, \ldots, m-1\}$, set $d_j := x_j \oplus x_0$ and form the $(m-1) \times n$ matrix D whose rows are $d_1, \ldots, d_{m-1}$. For any $x_0, x \in S$ we have $(x_0 \oplus x) \cdot u \equiv 0 \pmod 2$. Hence each d_j lies in the subspace

$$V = \{v \in \mathbb{F}_2^n \mid v \cdot u \equiv 0 \pmod 2\}, \qquad \dim V = n - 1.$$

Because the map $x \mapsto x_0 \oplus x$ is a bijection and the x_j are i.i.d. uniform in S, the d_j are i.i.d. uniform in V. Thus $\mathrm{rowspan}(D) \subseteq V$, so $\mathrm{rank}(D) \leq n - 1$. By the rank–nullity theorem,

$$\mathrm{rank}(D) + \mathrm{nullity}(D) = n.$$

If $\mathrm{rank}(D) = n - 1$, then $\mathrm{nullity}(D) = 1$, so $\ker(D)$ is the one-dimensional space spanned by u; consequently the unique nonzero solution to $Du' = 0$ is u. Gaussian elimination finds it in time polynomial in n and $\log(1/\epsilon)$. It remains to show that $\mathbb{P}(\mathrm{rank}(D) = n - 1) \geq 1 - \epsilon$. If the rows fail to span V, then all d_j lie in some proper $(n-2)$-dimensional subspace $W \subset V$. For any fixed W, we have

$$\mathbb{P}(d_j \in W) = \frac{|W|}{|V|} = \frac{2^{n-2}}{2^{n-1}} = \tfrac{1}{2},$$

independently across j, hence $\mathbb{P}(\forall j,\ d_j \in W) = 2^{-m+1}$. The number of $(n-2)$-dimensional subspaces of a $(n-1)$-dimensional $\mathbb{F}_2$-space is given by the Gaussian binomial

$$\binom{n-1}{n-2}_2 = \binom{n-1}{1}_2 = \frac{2^{n-1}-1}{2-1} = 2^{n-1}-1.$$

By a union bound,

$$\mathbb{P}(\operatorname{rank}(D) < n-1) \le (2^{n-1}-1)\,2^{-m+1} < 2^{n-m} \le \epsilon,$$

for our choice of m. □

Lower bound. Next, we will show that the bound is tight in n, i.e., $\Omega(n)$ samples are necessary. We follow a similar proof as in [2], and we find it convenient to talk about *success probability* p instead of expected score V, but these are related as $V = 2p - 1$.

Theorem 2. *Fix $n \ge 1$. Any classical algorithm that, given q samples from $S \in \mathcal{F}_{\mathrm{BV}}$, outputs some $y \in \bar{S}$ has success probability at most*

$$\frac{1}{2} + \frac{1}{2\left(2^{n-q+1}-1\right)} \qquad \textit{for } 1 \le q \le n.$$

Proof. Just as in [2], we assume query access to S by an oracle which on input i returns the i-th element of S, assuming arbitrary (and unknown) ordering of its elements. By Yao's minimax principle and since the index-query model only strengthens classical access relative to sampling [2], it suffices to bound the success of any deterministic index-query algorithm against the hard distribution where (u, b) is uniform over $\{0,1\}^n \setminus \{0\}^n \times \{0,1\}$ and $S = S_{u,b}$. Let $x_1, \ldots, x_q$ be any q distinct elements returned by the oracle. Let $W = \operatorname{span}\{x_1, \ldots, x_q\}$ with $r = \dim W \le q$. The constraints $x_i \in S$ mean that $u \cdot x_i = b$ for all i. The set of feasible pairs (u, b) consistent with the observed $x_1, \ldots, x_q$ is a linear subspace of $\mathbb{F}_2^n \times \mathbb{F}_2$ of dimension

$$d = n - r + 1,$$

since we have an $(n+1)$-dimensional space with r independent linear constraints. Under the uniform prior on (u, b) with $u \neq \{0\}^n$, conditioning on S containing $x_1, \ldots, x_q$ yields a uniform posterior over the remaining feasible pairs. This follows from a similar argument as in [2]: for any fixed subset $T \subseteq \{0,1\}^n$ of rank r, the number of pairs (u, b) with $T \subseteq S$ is the same. Their number is

$$N_{\mathrm{rem}} = 2^d - 1 = 2^{n-r+1} - 1.$$

Equivalently, the feasible u's form a d-dimensional linear subspace in $\mathbb{F}_2^n$. Write U for the associated d-dimensional linear subspace. Conditioned on the constraints being satisfied, u is uniform over $U \setminus \{0\}^n$ and b is fixed by $b \equiv u \cdot x_i \pmod 2$. Let the algorithm output any $y \notin \{x_1, \ldots, x_q\}$ such that there exists an $u' \in U \setminus \{0\}^n$ with $u' \cdot y \equiv 1 - b(u') \pmod 2$, where $b(u') \equiv u' \cdot x_1 \pmod 2$. We may assume without loss of generality that the algorithm chooses such a y, since any y with $u' \cdot y \equiv b(u') \pmod 2$ for all $u' \in U$ would make the success probability zero. For such a choice of y, the map $\phi_y : U \to \{0,1\}$ defined as

$$\phi_y(u') = u' \cdot (y + x_1) \pmod 2,$$

is a nonzero linear functional on the d-dimensional space U. Note that $\phi_y(u') = u' \cdot y + b(u') \equiv 1 \pmod 2$ if and only if $u' \cdot y \equiv 1 - b(u') \pmod 2$, which means that $\phi_y(u') = 1$ if and only if the algorithm succeeds for $(u', b(u'))$. Note that $\phi_y(u')$ takes each value in $\{0,1\}$ exactly 2^{d-1} times as u' ranges over U. Excluding $u' = \{0\}^n$ removes at most one element and never changes the

count for value 1, so among the $N_{\text{rem}} = 2^d - 1$ feasible inputs, exactly 2^{d-1} yield $\phi_y(u') = 1$. Therefore

$$\mathbb{P}(\text{success} \mid x_1, \ldots, x_q) \leq \frac{2^{d-1}}{2^d - 1},$$

where the success probability is with respect to the distribution over inputs. Since this is monotonically decreasing in d, this is minimized when $r = q$, i.e., the queried elements are linearly independent, giving $d = n - q + 1$. Since this holds for the constraints given any elements $x_1, \ldots, x_q$ and for any output y as specified above, the following final bound on the success probability holds:

$$\frac{2^{n-q}}{2^{n-q+1} - 1} = \frac{1}{2} + \frac{1}{2\big(2^{n-q+1} - 1\big)}. \tag{10}$$

□

Therefore, to achieve *constant* success probability $p > \frac{1}{2}$, we require $q = \Omega(n)$. Note that for $q = 1$, we recover Eq. (9) when we map Eq. (10) to the expected score via $V = 2p - 1$.

Supplementary Note 3: The rigidity of complement sampling

In this Supplementary Note we prove results for complement sampling which allows us to play the game without a trusted state preparation, and trusted strategy instead. We also prove that if both parties use untrusted devices then they could in principle collude to achieve perfect score. Recall the definition of rigidity score, $R(\mathcal{F}, h) = \mathbb{E}_S \mathbb{E}_c \mathbb{E}_{y|S,c}[(-1)^c \sigma(S, y)]$. We begin with a lower bound on the expected fidelity of state preparation when the player adopts the optimal quantum strategy Q.

Theorem 3. *Let n be a positive integer, let $\epsilon \in [0, 1]$, and let $\mathcal{F} = \{S_i | S_i \subseteq \{0,1\}^n, |S_i| = 2^{n-1}\}$ Suppose the referee has access to an untrusted quantum device that on input $S \in \mathcal{F}$ prepares an unknown state ρ_{SE}. The referee samples $S \in \mathcal{F}$ and $c \in \{0, 1\}$ uniformly at random: if $c = 0$, $\rho_S = \text{tr}_E(\rho_{SE})$ is sent to the player, who uses the optimal quantum strategy to return a n-bit string y; if $c = 1$ instead, the referee measures ρ_S in the computational basis to obtain a n-bit string y. Then*

$$R(\mathcal{F}, Q) \geq 1 - \epsilon \implies \mathbb{E}_S[\langle S| \rho_S |S\rangle] \geq 1 - \epsilon.$$

Proof. Fix the subset S and define the projector $\Pi_S \coloneqq \sum_{x \in S} |x\rangle\langle x|$. Then

$$\begin{aligned}
\mathbb{E}_c \mathbb{E}_{y|S,c}[(-1)^c \sigma(S, y)] &= \frac{1}{2}\mathbb{E}_{y|S,c=0}[\sigma(S, y)] + \frac{1}{2}\mathbb{E}_{y|S,c=1}[-\sigma(S, y)] \\
&= \frac{1}{2}\left(1 - 2\,\text{tr}\Big(\rho_S U \Pi_S U^\dagger\Big)\right) + \frac{1}{2}\left(2\,\text{tr}(\rho_S \Pi_S) - 1\right) \\
&= \text{tr}\Big(\rho_S (\Pi_S - U \Pi_S U^\dagger)\Big),
\end{aligned}$$

where we substitute the scores for when the referee measures directly and for when the player plays the optimal quantum strategy. Let us define the projector $\Pi_+ \coloneqq |+^n\rangle\langle+^n|$ so that under the optimal quantum strategy we have $U = 2\Pi_+ - I$. It is easy to verify the following properties:

$$\Pi_+ \Pi_S = \frac{1}{\sqrt{2}} |+^n\rangle\langle S|, \qquad \Pi_S \Pi_+ = \frac{1}{\sqrt{2}} |S\rangle\langle+^n|, \qquad \Pi_+ \Pi_S \Pi_+ = \frac{1}{2} |+^n\rangle\langle+^n|.$$

Thus

$$\begin{aligned}\Pi_S - U\Pi_S U^\dagger &= \Pi_S - (2\Pi_+ - I)\Pi_S(2\Pi_+ - I)\\ &= -4\Pi_+\Pi_S\Pi_+ + 2\Pi_S\Pi_+ + 2\Pi_+\Pi_S\\ &= -2\left|+^n\right\rangle\!\left\langle+^n\right| + \sqrt{2}\left|S\right\rangle\left\langle+^n\right| + \sqrt{2}\left|S\right\rangle\left\langle+^n\right|\\ &= -\left(\sqrt{2}\left|+^n\right\rangle - \left|S\right\rangle\right)\left(\sqrt{2}\left\langle+^n\right| - \left\langle S\right|\right) + \left|S\right\rangle\!\left\langle S\right|\\ &= -\left|\bar{S}\right\rangle\!\left\langle\bar{S}\right| + \left|S\right\rangle\!\left\langle S\right|,\end{aligned}$$

where we complete the square and use that $\left|\bar{S}\right\rangle = \sqrt{2}\left|+^n\right\rangle - \left|S\right\rangle$. Plugging this result back into the previous step we have

$$\mathbb{E}_c\mathbb{E}_{y|S,c}[(-1)^c\,\sigma(S,y)] = \operatorname{tr}(\rho_S\left|S\right\rangle\!\left\langle S\right|) - \operatorname{tr}\left(\rho_S\left|\bar{S}\right\rangle\!\left\langle\bar{S}\right|\right) \le \left\langle S\right|\rho_S\left|S\right\rangle. \tag{11}$$

Taking the expectation over subsets we have

$$\mathbb{E}_S\mathbb{E}_c\mathbb{E}_{y|S,c}[(-1)^c\,\sigma(S,y)] \le \mathbb{E}_S[\left\langle S\right|\rho_S\left|S\right\rangle],$$

and thus

$$R(\mathcal{F},Q) \ge 1-\epsilon \implies \mathbb{E}_S[\left\langle S\right|\rho_S\left|S\right\rangle] \ge 1-\epsilon.$$

□

Theorem 3 can be used to certify a lower bound on the fidelity of the referee's state preparation, assuming a trusted player device. Let us also look at the converse statement.

Proposition 1. *In the setting of Theorem 3, we also have that*

$$\mathbb{E}_S[\left\langle S\right|\rho_S\left|S\right\rangle] \ge 1-\epsilon \implies R(\mathcal{F},Q) \ge 1-2\epsilon.$$

Proof. Note that $\left|S\right\rangle\!\left\langle S\right| + \left|\bar{S}\right\rangle\!\left\langle\bar{S}\right| \preceq I$ because the two projectors are rank-1 and onto orthogonal subspaces. Taking the expectation of both sides and rearranging we have $\operatorname{tr}\left(\rho_S\left|\bar{S}\right\rangle\!\left\langle\bar{S}\right|\right) \le 1 - \operatorname{tr}(\rho_S\left|S\right\rangle\!\left\langle S\right|)$. Plugging this into the equality in Eq. (11) we have

$$\mathbb{E}_c\mathbb{E}_{y|S,c}[(-1)^c\,\sigma(S,y)] = \operatorname{tr}(\rho_S\left|S\right\rangle\!\left\langle S\right|) - \operatorname{tr}\left(\rho_S\left|\bar{S}\right\rangle\!\left\langle\bar{S}\right|\right) \ge 2\operatorname{tr}(\rho_S\left|S\right\rangle\!\left\langle S\right|) - 1.$$

Taking the expectation over subsets we have

$$\mathbb{E}_S\mathbb{E}_c\mathbb{E}_{y|S,c}[(-1)^c\,\sigma(S,y)] \ge 2\mathbb{E}_S[\left\langle S\right|\rho_S\left|S\right\rangle] - 1,$$

and it follows that

$$\mathbb{E}_S[\left\langle S\right|\rho_S\left|S\right\rangle] \ge 1-\epsilon \implies R(\mathcal{F},Q) \ge 1-2\epsilon.$$

□

Next, we upper-bound the rigidity score under the assumption that the player uses a specific trusted optimal classical strategy C. Specifically, upon receiving a string x, the player samples uniformly from $\{0,1\}^n \setminus \{x\}$, which is optimal in the noise-free setting because it is fixed-point-free (see Supplementary Note 1).

Proposition 2. *Let n be a positive integer, and let $\mathcal{F}_{\mathrm{BV}} = \{S_{u,b} \mid u \in \{0,1\}^n \setminus \{0^n\}, b \in \{0,1\}\}$. Suppose the referee has access to an untrusted device that on input $S \in \mathcal{F}_{\mathrm{BV}}$ outputs a random $x \in \{0,1\}^n$. Next, it samples $c \in \{0,1\}$ uniformly at random: if $c = 0$, x is sent to the player, who samples a uniform $y \in \{0,1\}^n \setminus \{x\}$ and returns it; if $c = 1$ instead, the referee checks if $x \in S$. Then*

$$R(\mathcal{F}_{\mathrm{BV}}, C) \le \frac{1}{2} + \frac{1}{2(2^n-1)}.$$

Proof. We first consider the $c = 0$ branch. Fix $S_{u,b} \in \mathcal{F}_{\mathrm{BV}}$ and a string $x \in \{0,1\}^n$ received by the player. If $x \in S_{u,b}$, we have

$$\mathbb{P}(y \notin S_{u,b} \mid S_{u,b}, x \in S_{u,b}) = \frac{2^{n-1}}{2^n - 1},$$

whereas, if $x \notin S_{u,b}$,

$$\mathbb{P}(y \notin S_{u,b} \mid S_{u,b}, x \in S_{u,b}) = \frac{2^{n-1} - 1}{2^n - 1}.$$

Using $\sigma(S, y) = 2 [\![y \notin S]\!] - 1$, it follows that

$$\mathbb{E}_{y|S_{u,b}}[\sigma(S_{u,b}, y)] = \begin{cases} \frac{1}{2^n-1}, & x \in S_{u,b}, \\ -\frac{1}{2^n-1}, & x \notin S_{u,b}. \end{cases}$$

In particular, for every $S_{u,b}$ and every possible received string x, this conditional expectation is at most $1/(2^n - 1)$ (which is the $x \notin S_{u,b}$ case). Hence, the expected score in the $c = 0$ branch is at most $1/(2^n - 1)$. In the $c = 1$ branch, the referee checks whether $x \in S_{u,b}$, which succeeds with probability at most 1. Since c is uniform, we obtain

$$R(\mathcal{F}_{\mathrm{BV}}, C) \leq \frac{1}{2} \cdot \frac{1}{2^n - 1} + \frac{1}{2} \cdot 1 = \frac{1}{2} + \frac{1}{2(2^n - 1)}.$$

□

Finally, we now briefly explain why the assumption that the classical strategy is trusted is necessary when the device is untrusted. If both the state preparation and the player's device are untrusted, then they can collude in a way that passes the relevant checks while producing samples that are far from uniform over the intended subset.

Proposition 3. *Let n be a positive integer. There exists a classical colluding strategy for an untrusted state preparation and player's device such that*

$$R(\mathcal{F}_{\mathrm{BV}}, C) = 1.$$

In this colluding strategy, for every $S_{u,b} = \{x \in \mathbb{F}_2^n \mid u \cdot x = b\}$, the state preparation and the player only consider a subset of $S_{u,b}$ of size 2^{n-2}.

Proof. The case of $n = 1$ is trivial since the state preparation can send the only x in the subset and the player can output $y = x \oplus 1$. Let us now consider $n \geq 2$. Choose a matrix A such that both A and $I - A$ are invertible. Such a matrix exists for every $n \geq 2$: identify $\mathbb{F}_2^n$ with the field $\mathbb{F}_{2^n}$ and choose $\alpha \in \mathbb{F}_{2^n} \setminus \{0, 1\}$, which is possible since $n \geq 2$. Let A be the $\mathbb{F}_2$-linear map given by multiplication by α. Then A is invertible because $\alpha \neq 0$, and $I - A$ is invertible because it is multiplication by $1 - \alpha \neq 0$. For each $S_{u,b}$, define

$$G_{u,b} = \{x \in \mathbb{F}_2^n \mid u \cdot x = b,\ u \cdot Ax = b \oplus 1\}.$$

We claim that $|G_{u,b}| = 2^{n-2}$. Indeed, the two constraints are linear constraints on $\mathbb{F}_2^n$. They are independent because the corresponding linear forms are u and $A^T u$, and $A^T u \neq u$: otherwise $(I - A^T)u = 0$, contradicting the invertibility of $I - A$. Hence the solution space has size 2^{n-2}. The colluding strategy is now simple. On input $S_{u,b}$, the state preparation produces x uniformly from $G_{u,b}$ and sends it to the player. By definition, $x \in S_{u,b}$. The player, upon receiving x, outputs $y = Ax$. Again by definition of $G_{u,b}$, we have $u \cdot Ax = b \oplus 1$, and therefore $y \notin S_{u,b}$. Thus the state preparation always passes the membership check and the player always outputs an element of the complement, which implies that $R(\mathcal{F}_{\mathrm{BV}}, C) = 1$. □

The only way to detect the above attack from samples is to distinguish the uniform distribution on $S_{u,b}$, which has support size 2^{n-1}, from the uniform distribution on the hidden subset $G_{u,b} \subseteq S_{u,b}$, which has support size 2^{n-2}. This requires exponentially many samples in the black-box sampling model. Hence a polynomially bounded referee cannot reliably rule out such a colluding strategy using statistics alone.

Supplementary Note 4: Analysis of the game with pseudorandom permutations subsets

As discussed in the main text, a subset state can be generated by applying a permutation P to $|0\rangle \otimes |+^{n-1}\rangle$. However, a reversible circuit implementation has exponential depth for general P. The original complement sampling game [2] proposes to use strong pseudorandom permutations (PRPs) instead. Strong PRPs typically have polynomial circuit depth and are indistinguishable from truly random permutations in polynomial time, even when given access to their inverse. The following result shows that complement sampling with these subsets remains hard in the average case for any classical strategy.

Theorem 4 (restatement of Thm. 7 in Ref. [2]). *For some $n \in \mathbb{N}$, let $P : \{0,1\}^n \times \{0,1\}^\lambda \to \{0,1\}^n$ be a strong pseudorandom permutation with security parameter $\lambda = n$. Given a key $k \in \{0,1\}^\lambda$, let S be the image of $\{0ik \mid i \in \{0,1\}^{n-1}\}$ under P, and $\bar{S}$ be the image of $\{1ik \mid i \in \{0,1\}^{n-1}\}$ under P. Let $O_{\text{index}} : [2^{n-1}] \to \{0,1\}^n$ be the oracle that on input i returns the i-th element in S. Picking a key uniformly at random, for all polynomial-time algorithms A that make a polynomial number of queries to O_{index} it must hold that*

$$\mathbb{P}\left(A^{O_{\text{index}}} \text{ outputs a } y \in \bar{S}\right) \leq \frac{1}{2} + \text{negl}(n).$$

Luby and Rackoff [3] showed that strong PRPs can be constructed from PRPs. PRPs used in cryptographic protocols such as Advanced Encryption Standard (AES) and Data Encryption Standard (DES) could be used for our purpose. However, they are made of rigid circuit constructions (the block ciphers) operating on a fixed number of bits (the block size), which would complicate the scaling analysis. To explore more flexible constructions we instead consider ensembles of ϵ-approximate k-wise independent PRPs. When applied to any fixed initial bit string, these PRPs produce distributions of bit strings that are indistinguishable from the uniform distribution up to the k-th moment and up to total variation distance ϵ.

Definition 1 (restatement of Def. 1 in Ref. [4]). *A distribution P on permutations of $\{0,1\}^n$ is said to be ϵ-approximate k-wise independent if for all distinct $x(1), \ldots, x(k) \in \{0,1\}^n$, the distribution of $(g(x(1)), \ldots, g(x(k)))$ for $g \sim P$ has total variation distance at most ϵ from the uniform distribution on k-tuples of distinct strings from $\{0,1\}^n$*

In the same paper, the authors provide circuit constructions based on layers of random reversible operations.

Theorem 5 (restatement of Thm. 3 in Ref. [4]). *A random one-dimensional brickwork circuit of three-bit gates and depth $n \cdot \tilde{\mathcal{O}}(k^2)$ computes a permutation that is 2^{-nk}-approximate k-wise independent.*

More generally, the idea is to build a reversible circuit from layers of random three-bit operations, so that inputs and outputs become increasingly uncorrelated as the number of layers grows. To explore variations of this construction, we define three circuit parameters: layer type, operation type, and number of layers. The *layer type* prescribes the location of the operations. We consider either placing them in a one-dimensional nearest-neighbors brickwork (NN), or as a layout of random triplet (RT). More specifically, NN is a sequence of $\frac{n}{3}$ three-bits operations starting from the first qubit $(1,2,3),(4,5,6),\ldots,(n-2,n-1,n)$, followed by another sequence starting from the second qubit $(2,3,4),(5,6,7),\ldots,(n-1,n,1)$. In RT we proceed similarly except that each triplet is selected at random. For example, we could have $(1,6,8),(5,2,9),(3,4,7)$, followed by $(3,7,8),(1,5,6),(9,2,4)$. The RT construction is motivated by the full-connectivity of Quantinuum H-series devices and is similar in

Layer type	Operation type	Number of layers	Average 2-qubit gates	Worst-case 2-qubit gates
NN	DES2	1	$2n$	$10n/3$
NN	DES2	n	$2n^2$	$10n^2/3$
NN	S8	1	$20n/3$	$12n$
NN	S8	n	$20n^2/3$	$12n^2$
RT	DES2	1	$2n$	$10n/3$
RT	DES2	n	$2n^2$	$10n^2/3$
RT	S8	1	$20n/3$	$12n$
RT	S8	n	$20n^2/3$	$12n^2$

Supplementary Table 1: *Pseudorandom permutations constructions.*

spirit to the construction used for random circuit sampling in Ref. [5]. The *operation type* can be any gate set which is universal for reversible classical computation. We consider DES2 gates [6], which are also used in the proof of Theorem 5. DES2 gates act on the three bits as $(b_1, b_2, b_3) \to g(b_1, b_2, b_3 \oplus f(b_1, b_2))$ for any of the $2^{2^2} = 16$ possible binary functions f, and any of the $3! = 6$ possible reordering g of the output bits. Thus there are 96 possible instantiations of a DES2 gate on three bits. In addition, we consider S8 gates that act as arbitrary permutations of the eight basis states of three bits. S8 is a superset of DES2 and there exists $8! = 40320$ possible instantiations on three bits. Finally, we consider either a constant *number of layers* $L = 1$, or a number of layers that scales with the number of bits $L = n$. We will refer to our PRP constructions using the name convention: layer type - operation type - number of layers. For example, NN-DES2-n consists of n layers of nearest-neighbors DES2 gates.

Supplementary Table 1 shows the constructions considered in this study and reports the average- and worst-case scaling of number of two-qubit gates. These are estimated by recalling that each layer contains $2n/3$ operations and by decomposing each operation to quantum gates as explained in Supplementary Note 5. The estimates are validated in Supplementary Fig. 1(b) where we take n to be a multiple of 3 and report the average number of two-qubit gates of 100 random instances. Each circuit is compiled for the Quantinuum H-series native gate set $R_{ZZ}(\theta) = e^{-i\frac{\theta}{2}Z\otimes Z}$ and $U(\theta, \phi) = e^{-i\frac{\theta}{2}(\cos(\phi)X+\sin(\phi)Y)}$ using optimization level 3 of the `pytket` software package [7]. The number of gates increases linearly for single layer constructions and quadratically for n layer constructions, as expected. We note a significant overhead for implementing S8 (blue lines) compared to DES2 gates (red lines), while the layer type did not have a significant impact.

Next we numerically verify that our PRP constructions are ϵ-approximate k-wise independent where, following Theorem 5, we choose $\epsilon = 2^{-kn}$. The idea is to sample random instances, apply them to the all-zero bit string, and measure in the computational basis. We then compare the distribution of output bit strings p_i to the uniform distribution in terms of total variation distance $\text{TVD} = \frac{1}{2}\sum_{i=1}^{2^n}|p_i - \frac{1}{2^n}|$. A sample-optimal algorithm for this task is given in Ref. [8]. It computes the random variable $S = \frac{1}{2}\sum_{i=1}^{2^n}|\tilde{p}_i - \frac{1}{2^n}|$ where $\tilde{p}_i$ are empirical probabilities estimated from the M samples. If S is larger than a specific threshold, the algorithm outputs 'NO', otherwise it returns 'YES'. The authors prove the following theorem.

Theorem 6 (restatement of Thm. 3 in Ref. [8])**.** *Let p be an unknown distribution over 2^n elements. There exists a universal constant $C > 0$ such that the following holds. Given $M \geq \frac{C}{\epsilon^2}(\sqrt{2^n \log\frac{1}{\delta}} + \log\frac{1}{\delta})$ samples from p, there exists an algorithm that outputs 'YES' with probability $1 - \delta$ if p is the uniform distribution, and 'NO' with probability $1 - \delta$ if the total variation distance to the uniform distribution is $\geq \epsilon$.*

Using this theorem with $\epsilon = 2^{-kn}$ and $k = 1$, we observe that $\mathcal{O}(2^{2.5n})$ samples are sufficient for the algorithm in [8] to solve the decision problem. In our work we use S as an estimate for the TVD, so we don't threshold its value. For each PRP construction in Table 1 we sample $10n \times 2^{2.5}$ random instances, where the factor of $10n$ is chosen as the maximum we can achieve with our

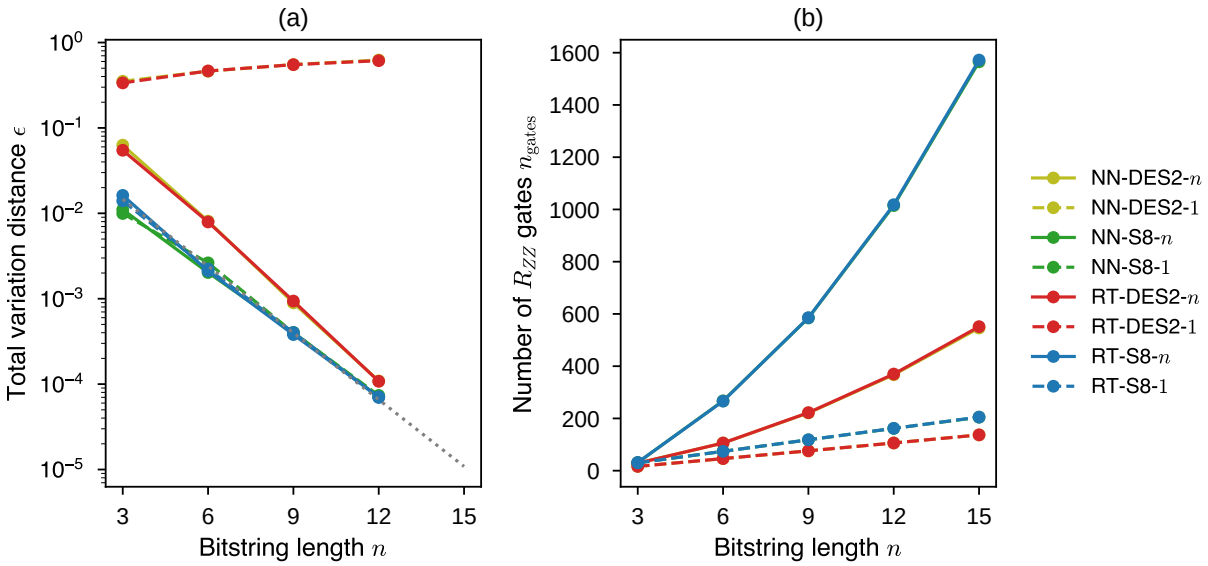

Supplementary Figure 1: *Numerical analysis of pseudorandom permutations constructions.* (a) Total variation distance ϵ to the uniform distribution versus number of bits n. (b) Number of two-qubit R_{ZZ} gates n_{gates} required for a reversible circuit implementation optimized for Quantinuum H2-1. Each point is averaged over 100 random instances. Data corresponding to NN constructions lie directly below the corresponding RT constructions.

limited computational capabilities. Note that we could not analyze the second moment, $k = 2$, because it requires the concatenation of 2 input bit strings (see Def. 1) leading to a number of samples scaling as $\mathcal{O}(2^{5n})$.

Supplementary Fig. 1(a) shows that ϵ decreases exponentially in n for all constructions except NN-DES2-1 and RT-DES2-1. This indicates that a single layer of DES2 gates does not result in sufficient information scrambling, while a single layer of S8 gates does. With regard to the connectivity of the layers, we do not observe a significant difference between nearest-neighbors (NN) and random triplets (RT). Finally, from [4, Theorem 3] it follows that NN-DES2-n asymptotically provides 2^{-n}-approximate 1-wise independent pseudorandom permutations, which is confirmed here numerically (yellow solid line). We extrapolate to $n = 15$ (dotted gray line) for RT-S8-1 and $n = 36$ (not shown) which are setups used later on in hardware experiments. Finally, we shall point out a subtlety: since S8 gates are *exact* permutations on three bits, we would expect the TVD for $n = 3$ (green and blue lines) to be zero. This does not happen because we use a finite number of samples, albeit sufficient for our purposes due to Theorem 6. We confirm that when using ten times more samples the TVD of S8 decreases, and that of DES2 remains the same. Since our analysis concerns upper bounds to the TVD, this fact does not affect our conclusions.

We use the data in Supplementary Fig. 1 to choose suitable settings for executing the complement sampling game with PRP subsets on Quantinuum H-series hardware. To this end, we shall take into account also the linear cost of two-qubit gates for implementing both teleportation and swapper circuits (see Supplementary Note 5). We choose to execute RT-DES2-n with $n = 12$, yielding $n_{\text{qubits}} = 41$ and an average of $n_{\text{gates}} = 370$, as well as RT-S8-1 with $n = 15$, yielding $n_{\text{qubits}} = 51$ and an average of $n_{\text{gates}} = 205$. We also execute the game without the teleportation protocol, and with the referee and player sharing the same register. This allows us to explore much longer bit strings. We consider RT-DES2-n with $n = 36$, yielding $n_{\text{qubits}} = 53$ and an average of $n_{\text{gates}} = 438$.

We compare the empirical score $\hat{V}$ of the implemented quantum strategy against the score of an optimal classical strategy that can somehow exploit the approximation error in the PRP

n	PRP	ϵ	teleportation	n_{qubits}	n_{gates} avg	n_{gates} std
12	RT-DES2-n	1.078×10^{-4}	yes	41	370.47	19.96
15	RT-S8-1	1.087×10^{-5}	yes	51	205.16	9.45
36	RT-S8-1	3.639×10^{-11}	no	53	438.10	14.72

Supplementary Table 2: *Summary of the experimental settings for PRP subset states.*

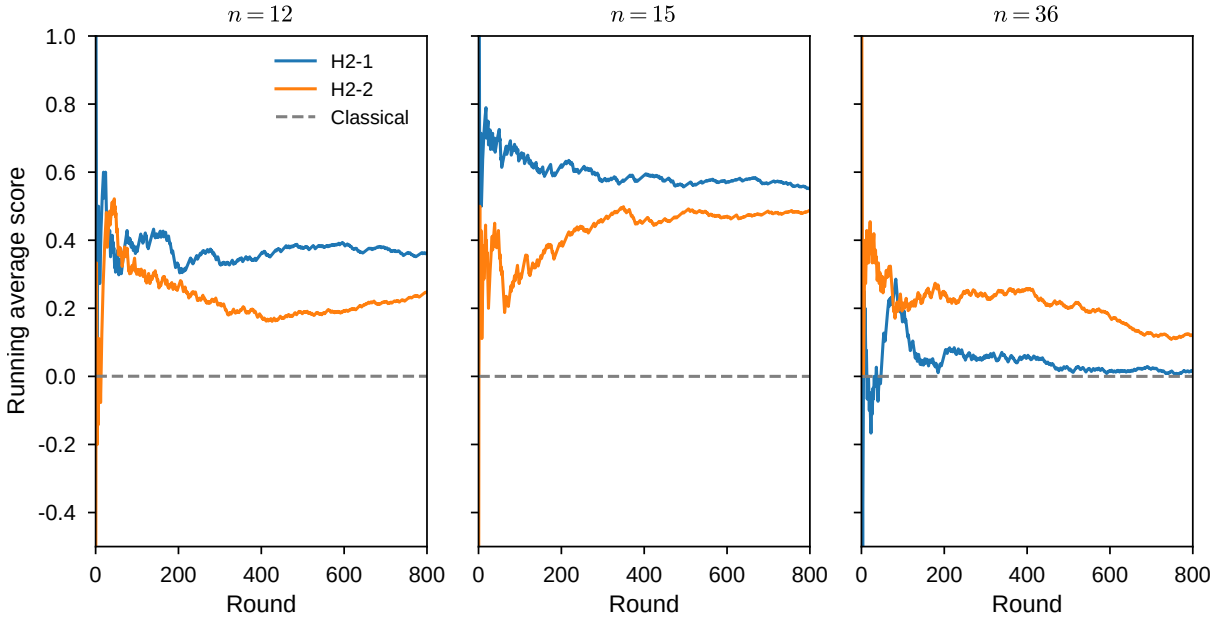

Supplementary Figure 2: *Experimental results for PRP subset states.* Running average of the player's score plotted against the number of rounds. The optimal classical strategy achieves the score indicated by the gray line.

family. Denoting the family of approximate PRPs as $\mathcal{F}_{\text{PRP},\epsilon}$, the optimal classical strategy achieves $V(\mathcal{F}_{\text{PRP},\epsilon}, C) = \frac{1}{2^n-1} + 2\epsilon$. Here ϵ is found by extrapolating the data in Supplementary Fig. 1 and a factor of 2 arises from the conversion from TVD to score. Each game is executed for 800 rounds on the H2-1 and H2-2 quantum computers. This amounts to running 4800 quantum circuits and obtaining a single sample from each. Our setup is summarized in Supplementary Table 2. At the time of these experiments in July 2025, H2-1 had two-qubit gate infidelity of $(1.05 \pm 0.08) \times 10^{-3}$ and memory error per qubit and depth-1 circuit time of $(2.03 \pm 0.23) \times 10^{-4}$. H2-2 had two-qubit infidelity of $(1.10 \pm 0.17) \times 10^{-3}$ and memory error per qubit and depth-1 circuit time of $(2.26 \pm 0.29) \times 10^{-4}$. These were measured via randomized benchmarking of all qubits [9–12].

Supplementary Fig. 2 shows the running average of the empirical score during the game. In almost all settings the quantum strategy outperforms the optimal classical strategy (dashed gray line) significantly. This provides evidence of a large violation of classicality, albeit conditional on the existence of PRPs or, equivalently, on the existence of one-way functions [13]. Noting the low performance on the $n = 36$ run with H2-1, we choose to employ H2-2 in the experiments presented in the main text.

Supplementary Note 5: Details about the implementation

Quantum teleportation. A quantum communication channel is used in some experiments with both BV and PRP subset states. A simple way to simulate a quantum communication channel between two parties is to swap their qubits. This 'symmetric' communication channel implies that while the referee sends a state to the player, the player simultaneously sends back

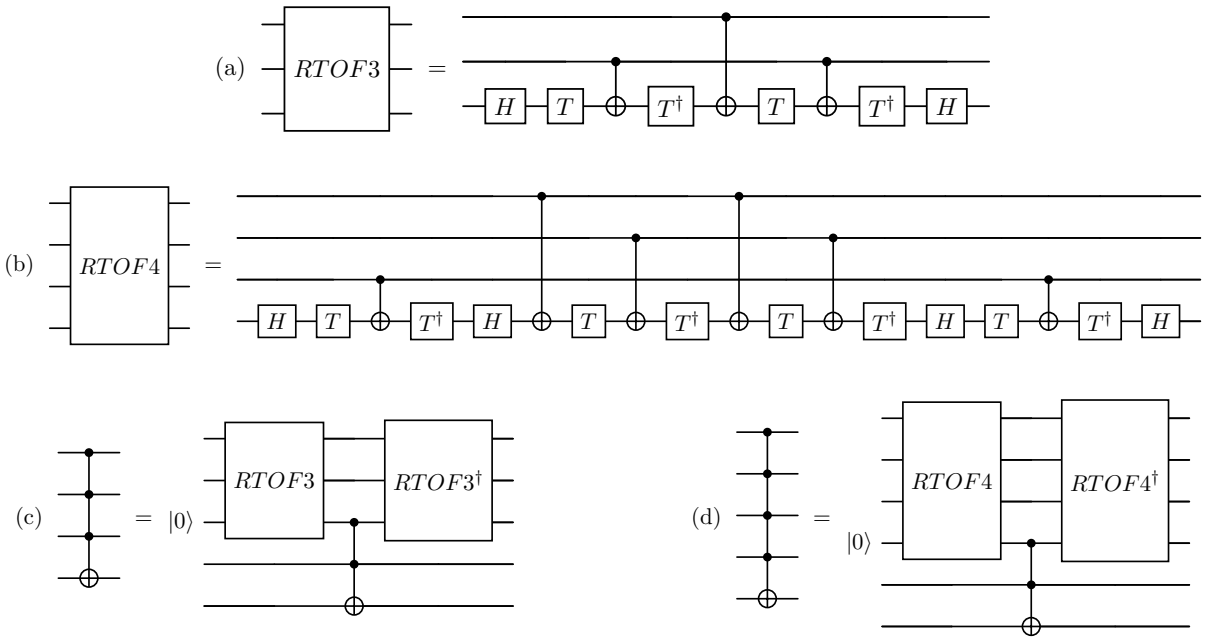

Supplementary Figure 3: (a) Three-qubit Toffoli gate implemented up to a relative phase and with optimal CX count. (b) Four-qubit Toffoli gate implemented up to a relative phase. (c) Four-qubit Toffoli and (d) five-qubit Toffoli implemented using a single ancilla qubit. RTOF3 is self-adjoint, but we keep the † symbol as a guidance for the recursive construction of high-order Toffoli gates explained in the text. All circuits are taken from Ref. [15].

a state to the referee. We opted for a more realistic and resource consuming 'asymmetric' quantum communication channel. This is provided by the well-known quantum teleportation protocol [14]. The two parties share a set of Bell pairs, the referee performs local operations and measurements on them, and then communicates the outcomes via a classical channel. The player applies local operations on the Bell pairs accordingly, effectively recovering the referee's state.

In the main text, Fig. 1C shows the textbook implementation of quantum teleportation for a single qubit in the state $|S\rangle$. From Supplementary Fig. 4, a CX can be decomposed into a R_{ZZ} gate up to a few single-qubit gates. The mid-circuit measurements are followed by classical communication and correcting operations. These are implemented on the Quantinuum H2 hardware using conditional gates. For the implementation of an n-qubit state we simply apply this circuit to each qubit and associated ancilla.

In summary, for a state of n qubits the communication between referee and player requires n ancilla qubits, $2n$ two-qubit gates, $2n$ mid-circuit measurements, and $2n$ conditional single-qubit gates.

Compilation of the swapper circuit. The swapper circuit is used in all our experiments. For subsets of size 2^{n-1} the optimal swapper is the n-qubit Grover diffusion operator, up to a phase. In the main text, Fig. 1D, right panel, shows the high level circuit for $n = 4$. Clearly the main implementation bottleneck is the multi-controlled Toffoli gate. We compiled the circuit to the native gate set of Quantinuum H-series for several values of n using `pytket`. We found that the number of two-qubit gates almost exactly matches the curve $4n^2 - 4n - 0.5$. However, in our game the player is allowed to use ancilla qubits as a workspace for a more efficient compilation. While there exist a plethora of techniques that could be used to compile multi-controlled Toffoli gates, we choose to implement the method in Ref. [15]. In particular we utilize the following result.

$f(b_1, b_2)$	decomposition	two-qubit gates
0	I	0
$b_1 \wedge b_2$	$\mathrm{CCX}(b_1, b_2, b_3)$	5
$b_1 \wedge \neg b_2$	$\mathrm{X}(b_2)\mathrm{CCX}(b_1, b_2, b_3)\mathrm{X}(b_2)$	5
b_1	$\mathrm{CX}(b_1, b_3)$	1
$\neg b_1 \wedge b_2$	$\mathrm{X}(b_1)\mathrm{CCX}(b_1, b_2, b_3)\mathrm{X}(b_1)$	5
b_2	$\mathrm{CX}(b_2, b_3)$	1
$b_1 \oplus b_2$	$\mathrm{CX}(b_1, b_3)\mathrm{CX}(b_2, b_3)$	2
$b_1 \vee b_2$	$\mathrm{X}(b_1)\mathrm{X}(b_2)\mathrm{CCX}(b_1, b_2, b_3)\mathrm{X}(b_3)\mathrm{X}(b_2)\mathrm{X}(b_1)$	5
$\neg(b_1 \vee b_2)$	$\mathrm{X}(b_1)\mathrm{X}(b_2)\mathrm{CCX}(b_1, b_2, b_3)\mathrm{X}(b_2)\mathrm{X}(b_1)$	5
$\neg(b_1 \oplus b_2)$	$\mathrm{CX}(b_1, b_3)\mathrm{CX}(b_2, b_3)\mathrm{X}(b_3)$	2
$\neg b_2$	$\mathrm{X}(b_2)\mathrm{CX}(b_2, b_3)\mathrm{X}(b_2)$	1
$b_1 \vee \neg b_2$	$\mathrm{X}(b_1)\mathrm{CCX}(b_1, b_2, b_3)\mathrm{X}(b_3)\mathrm{X}(b_1)$	5
$\neg b_1$	$\mathrm{X}(b_1)\mathrm{CX}(b_1, b_3)\mathrm{X}(b_1)$	1
$\neg b_1 \vee b_2$	$\mathrm{X}(b_2)\mathrm{CCX}(b_1, b_2, b_3)\mathrm{X}(b_3)\mathrm{X}(b_2)$	5
$\neg(b_1 \wedge b_2)$	$\mathrm{CCX}(b_1, b_2, b_3)\mathrm{X}(b_3)$	5
1	$\mathrm{X}(b_3)$	0

Supplementary Table 3: Decomposition of DES2 gates to quantum circuits in the gate set X, CX, and CCX.

Theorem 7 (restatement of Prop. 4 in Ref. [15])**.** *A size $n \geq 4$ multi-controlled Toffoli gate can be implemented by a circuit that uses $\lceil \frac{n-3}{2} \rceil$ ancillary qubits initialized in and returned to $|0\rangle$, $8n - 17$ T gates, $6n - 12$ CX gates, and $4n - 10$ Hadamard gates.*

The author gives a constructive proof which uses as building blocks the three-qubit Toffoli (TOF3), the three-qubit Toffoli up to a relative phase (RTOF3) and the four-qubit Toffoli up to a relative phase (RTOF4) illustrated in Supplementary Fig. 3(a) and (b). In Supplementary Fig. 3(c) we show the implementation of the four-qubit Toffoli (TOF4) using a single ancilla qubit. The top two qubits are used as controls for the RTOF3 with the ancilla as target. We then apply a TOF3 and finally we undo the RTOF3. To go from TOF4 to TOF5, Supplementary Fig. 3(d), we simply exchange the RTOF3s to RTOF4s. For a six-qubit Toffoli (not shown) we add an ancilla and exchange the TOF3 in the middle with a TOF4. As we increase the number of qubits the recursion rule is the following: i) when going from an odd number of controls to even, exchange the outermost RTOF3s to RTOF4s, and ii) when going from an even number of controls to odd, add an ancilla qubit and exchange the TOF3 in the middle to a TOF4.

Compilation of DES2 operations. DES2 operations are used in our experiments with PRPs. DES2 operations map any 3-bits computational basis state $|b_1, b_2, b_3\rangle$ to $|b_1, b_2, b_3 \oplus f(b_1, b_2)\rangle$ up to arbitrary order of the output bits. The reordering is implemented by at most two SWAP gates. On the H2 devices, the SWAP gate is handled by relabeling and transporting qubits, therefore not incurring any extra cost on the two-qubit gate count [9]. The function is defined as $f : \{0, 1\}^2 \to \{0, 1\}$, so there are $2^{2^2} = 16$ of them. Supplementary Table 3 shows all possible functions along with a decomposition in terms of X, CX and CCX gates. It is known that a CCX, also known as Toffoli gate, can be decomposed into five generic two-qubit gates [16] and that this number is necessary [17]. For completeness we report a possible two-qubits gate decomposition using the native gate set of Quantinuum H-series hardware in Supplementary Fig. 4. We conclude that DES2 operations require five two-qubit gates in the worst case, and three in average.

Compilation of S8 operations. S8 operations are only used in our experiments with PRPs. S8 operations permute the computational basis states of 3 bits. We approach the problem of compiling these operations numerically by brute-force. We use the `pytket` software package [7]

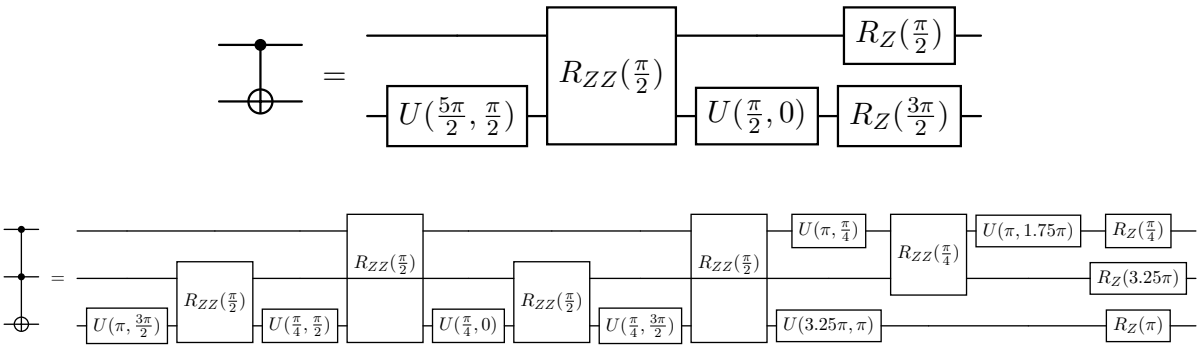

Supplementary Figure 4: Decomposition of CX (top) and CCX (bottom) gates using the native gate set of the Quantinuum H-series hardware. Here $R_{ZZ}(\theta) = e^{-i\frac{\theta}{2}Z\otimes Z}$, $R_Z(\theta) = e^{-i\frac{\theta}{2}Z}$ and $U(\theta, \phi) = e^{-i\frac{\theta}{2}(\cos(\phi)X + \sin(\phi)Y)}$.

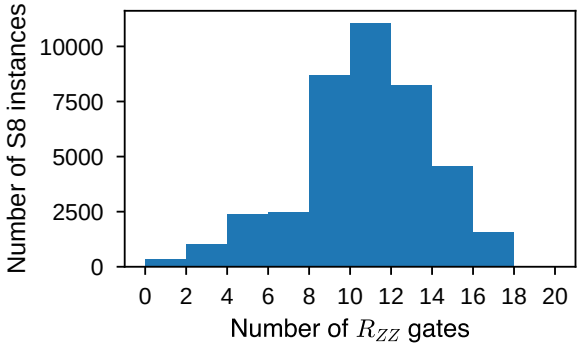

Supplementary Figure 5: Histogram of S8 instances that require a certain number of $R_{ZZ}(\theta)$ two-qubit gates.

and compile for the native Quantinuum H-series gate set. For each of the $8! = 40320$ permutations we construct the corresponding unitary using the `ToffoliBox` tool in `pytket`. We then compile each unitary to a quantum circuit and optimize it with the highest level of optimization (three), which consists of a collection of gate reduction techniques. Supplementary Fig. 5 shows the histogram of how many permutations yielded a certain number of two-qubit $R_{ZZ}(\theta)$ gates. We find that S8 operations require at most 18 two-qubit gates, and 10 on average.

Code's functionality. In [18] we provide access to the package that implements the circuits described in the previous sections. The user can run the example notebooks to generate the full circuits for the desired experiment and then inspect the results. For the hardware experiments, the authors ran the circuits on Quantinuum System Model H2 quantum computers via Quantinuum Nexus [19], but other backends could be selected to test such experiments.

Supplementary References

[1] H. Buhrman, O. Regev, G. Scarpa, and R. d. Wolf, "Near-Optimal and Explicit Bell Inequality Violations," *Theory of Computing* **8**, no. 27, 623–645 (2012).

[2] M. Benedetti, H. Buhrman, and J. Weggemans, "Provable and verifiable quantum advantage in sample complexity," *Phys. Rev. Lett.* **136**, 040601 (Jan, 2026).

[3] M. Luby and C. Rackoff, "How to construct pseudorandom permutations from pseudorandom functions," *SIAM Journal on Computing* **17**, no. 2, 373–386 (1988).

[4] W. Gay, W. He, N. Kocurek, and R. O'Donnell, "Pseudorandomness properties of random reversible circuits," in *Advances in Cryptology – CRYPTO 2025*, pp. 651–678. Springer Nature Switzerland, Cham, 2025.

[5] M. DeCross, R. Haghshenas, *et al.*, "Computational power of random quantum circuits in arbitrary geometries," *Phys. Rev. X* **15**, 021052 (2025).

[6] S. Even and O. Goldreich, "DES-like functions can generate the alternating group," *IEEE Transactions on Information Theory* **29**, no. 6, 863–865 (1983).

[7] S. Sivarajah, S. Dilkes, A. Cowtan, W. Simmons, A. Edgington, and R. Duncan, "t|ket⟩: a retargetable compiler for NISQ devices," *Quantum Science and Technology* **6**, no. 1, 014003 (2020).

[8] I. Diakonikolas, T. Gouleakis, J. Peebles, and E. Price, "Optimal identity testing with high probability," `arXiv:1708.02728` (2017).

[9] S. A. Moses, C. H. Baldwin, *et al.*, "A race-track trapped-ion quantum processor," *Phys. Rev. X* **13**, 041052 (2023).

[10] E. Magesan, J. M. Gambetta, and J. Emerson, "Scalable and robust randomized benchmarking of quantum processes," *Phys. Rev. Lett.* **106**, 180504 (2011).

[11] R. Harper, I. Hincks, C. Ferrie, S. T. Flammia, and J. J. Wallman, "Statistical analysis of randomized benchmarking," *Phys. Rev. A* **99**, 052350 (2019).

[12] Quantinuum, 2025. `https://docs.quantinuum.com/systems/user_guide/hardware_user_guide/performance_validation.html` [Accessed: 2025-10-16].

[13] J. Katz and Y. Lindell, *Introduction to Modern Cryptography*. Chapman & Hall/CRC Cryptography and Network Security Series. CRC Press, 2020.

[14] C. H. Bennett, G. Brassard, C. Crépeau, R. Jozsa, A. Peres, and W. K. Wootters, "Teleporting an unknown quantum state via dual classical and Einstein-Podolsky-Rosen channels," *Phys. Rev. Lett.* **70**, 1895–1899 (1993).

[15] D. Maslov, "Advantages of using relative-phase Toffoli gates with an application to multiple control Toffoli optimization," *Physical Review A* **93**, no. 2, 022311 (2016).

[16] D. P. DiVincenzo, "Quantum gates and circuits," *Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences* **454**, no. 1969, 261–276 (1998).

[17] N. Yu, R. Duan, and M. Ying, "Five two-qubit gates are necessary for implementing the Toffoli gate," *Phys. Rev. A* **88**, 010304 (2013).

[18] M. Rosenkranz, G. Marin-Sanchez, and M. Benedetti, "Supporting material for 'Unconditional and exponentially large violation of classicality'." 2026. Zenodo. `https://doi.org/10.5281/zenodo.17593255`.

[19] "Quantinuum Nexus." 2024. `https://nexus.quantinuum.com/`.