

ARTICLE

Received 17 Apr 2013 | Accepted 10 Nov 2013 | Published 5 Dec 2013

DOI: 10.1038/ncomms3915

Quantum replication at the Heisenberg limit

Giulio Chiribella¹, Yuxiang Yang¹ & Andrew Chi-Chih Yao²

No process in nature can perfectly clone an arbitrary quantum state. But is it possible to engineer processes that replicate quantum information with vanishingly small error? Here we demonstrate the possibility of probabilistic super-replication phenomena where N equally prepared quantum clocks are transformed into a much larger number of M nearly perfect replicas, with an error that rapidly vanishes whenever M is small compared with N^2 . The quadratic replication rate is the ultimate limit imposed by quantum mechanics to the proliferation of information and is fundamentally linked with the Heisenberg limit of quantum metrology.

¹Center for Quantum Information, Institute for Interdisciplinary Information Sciences, Tsinghua University, Beijing 100084, China. ²Institute for Interdisciplinary Information Sciences, Tsinghua University, Beijing 100084, China. Correspondence and requests for materials should be addressed to G.C. (email: gchiribella@mail.tsinghua.edu.cn).

No physical process can copy arbitrary quantum states on demand^{1,2}; if such a process existed, we could build a device that distinguishes quantum states with arbitrary precision, violating the uncertainty principle and enabling faster-than-light communication³. Probabilistic processes like stimulated emission, however, seem to evade this restriction. In Einstein's treatment of stimulated emission⁴, an excited atom interacting with a polarized photon is expected to produce sometimes a second photon with the same polarization, effectively delivering a perfect clone. True that some other times the atom will spontaneously emit a photon of random polarization, but still, when stimulated emission occurs, a perfect clone has been produced. If we had on our side a quantum version of Maxwell demon, who separates the photons produced by stimulated emission from those produced by spontaneous emission, we would be able to generate any desired number of clones with a non-zero probability. Unfortunately, our imaginary helper is not allowed by the laws of quantum mechanics, even with tiny probability, knowing whether or not stimulated emission took place would lead to a violation of the no-signalling principle (incidentally, this observation should be taken as a reminder that Einstein's treatment is just an approximation, whereas in the actual quantum dynamics stimulated and spontaneous emission happen in a coherent superposition). However, nothing forbids that other probabilistic processes, akin to spontaneous emission, could proliferate quantum information beyond any previously conceived limit. This possibility raises new fundamental questions: is it possible to engineer a process that duplicates a beam of N equally prepared particles, producing a beam of $2N$ almost perfect clones? What is the ultimate rate at which quantum information proliferate without significant errors?

Here we answer both questions: although it is impossible to duplicate arbitrary quantum states, we devise a probabilistic mechanism that transforms an input beam of N particles, equally prepared in a state $|\psi_t\rangle = e^{-iHt}|\psi\rangle$ generated by time evolution, into an output beam of magnified intensity, consisting of an overwhelming number M of nearly perfect clones with a small error that vanishes rapidly whenever M is small compared with N^2 . We name this new phenomenon super-replication and show that it is intrinsically probabilistic, by proving that deterministic processes can only produce a negligible number of nearly perfect replicas. For example, for 100 linearly polarized photons, super-replication allows one to produce 1,000 replicas with fidelity 99.9%, whereas the best deterministic process can only achieve fidelity 57%. In addition, we show that no physical process, deterministic or not, can proliferate quantum information at a rate larger than quadratic; any attempt to replicate quantum information beyond this limit is doomed to produce a joint output that has vanishing fidelity with the desired state. To explain the roots of this fundamental limitation, we establish a deep link between quantum cloning and the precision limits of quantum metrology^{5–11}, showing that the Heisenberg limit (HL) sets the ultimate bound to the replication rate of probabilistic processes, while the standard quantum limit (SQL) sets the corresponding bound in the deterministic regime.

Results

The SQL for information replication. Optimal cloning is a fundamental primitive in quantum information^{12–14}. Its goal is to transform N input copies of a quantum state $|\psi_x\rangle$, randomly drawn from a set $\{|\psi_x\rangle\}_{x \in X}$, into M approximate copies that are as faithful as possible. The simplest figure of merit here is the fidelity $F_{N \rightarrow M}$ between the M -particle state produced by the process and M exact copies of the desired state, evaluated on the worst-case input state. Inspired by information theory, we now

consider a sequence of cloning processes that transform N input copies into $M = M(N)$ approximate copies and we say that the replication is reliable if the replication error vanishes in the limit of large N . As an error measure, we use the infidelity $e_{N \rightarrow M} = 1 - F_{N \rightarrow M}$. Here the crucial question is how large can M grow as a function of N in a reliable replication process? To answer the question, we introduce the notion of replication rate, saying that a replication process has rate α if the number of extra copies scales like N^α . We say that a rate is achievable if there exists a sequence of reliable replication processes with that rate. We will now see that the achievable replication rates are determined by the precision limits of quantum metrology. Our first key result in this direction is a SQL for quantum replication: No deterministic process can reliably replicate a continuous set of quantum states at a rate larger than 1. In other words, deterministic processes can only embezzle from nature a negligible number of extra copies.

The derivation of the SQL for quantum replication is based on the SQL for quantum metrology¹⁵, applied to an arbitrary curve of states $\{|\psi_t\rangle | t \in (-\epsilon, \epsilon)\}$ contained in the set of states that we are trying to clone. A sketch of proof is as follows: the SQL states that the variance in the estimation of t from N copies is lower bounded by $V_{N,t} \geq \frac{c}{N}$, for a constant c that can be set to $c = 1$ with a suitable choice of parametrization. Now, suppose that there is a sequence of reliable deterministic processes with rate α . The N th process will produce an M -particle output state $\rho_{M,t}^{\text{out}}$, where $M \geq N + aN^\alpha$ for some constant $a > 0$, approaching the ideal target $|\psi_t\rangle^{\otimes M}$ in the large N limit. For simplicity, let us assume that the trace distance between $\rho_{M,t}^{\text{out}}$ and $|\psi_t\rangle^{\otimes M}$ vanishes as $O(M^{-\beta})$ for some exponent $\beta \geq 1$ (for sets of states of the form $\{|\psi_t\rangle = e^{-iHt}|\psi\rangle | t \in \mathbb{R}\}$, this assumption is lifted in Supplementary Note 1). Now, if two states are close, so is the variance in the estimation of t from these two states: for every fixed estimation strategy, we have the bound $V_{M,t}^{\text{out}} \leq V_{M,t} + \gamma M^{-\beta}$, where $V_{M,t}^{\text{out}}$ is the variance in the estimation of t from $\rho_{M,t}^{\text{out}}$, $V_{M,t}$ is the variance in the estimation of t from M copies and $\gamma > 0$ is a suitable constant. However, applying a deterministic transformation cannot reduce the variance of the optimal estimation strategy, denoted by $V_{N,t}^{\text{min}}$. Hence, we have the bound $V_{N,t}^{\text{min}} \leq V_{M,t}^{\text{out}} \leq V_{M,t} + \gamma M^{-\beta}$. Choosing the best estimation strategy, we then obtain $V_{N,t}^{\text{min}} \leq V_{M,t}^{\text{min}} + \gamma M^{-\beta}$. Now, since the SQL is asymptotically achievable by a suitable strategy (cf. Methods), for large N the bound becomes $1/N \leq 1/(N + aN^\alpha) + \gamma(N + aN^\alpha)^{-\beta}$, up to terms that are negligible with respect to N^{-1} . Clearly, this implies $\alpha \leq 1$ because otherwise we would have a contradiction for large N . This establishes the SQL for quantum replication. As a by-product of the derivation, we also have that no deterministic process can replicate information with error vanishing faster than N^{-4} . Indeed, by Taylor expanding on the right hand side of the inequality $1/N \leq 1/(N + aN^\alpha) + \gamma(N + aN^\alpha)^{-\beta}$, we obtain the condition $\beta \leq 2 - \alpha$, which implies $\beta \leq 2$. Using the relation between fidelity and trace distance¹⁶, this means that the replication error is lower bounded as $e_{N \rightarrow M} \geq O(N^{-4})$. In other words, for a reliable replication process the error cannot vanish faster than a low-degree polynomial.

The SQL limits not only deterministic replication but also some instances of probabilistic replication; for example, probabilistic cloning has no advantage for arbitrary pure states (this fact was observed in ref. 17 for 1-to- M cloning, but the argument can be easily extended to N -to- M cloning). More generally, it is easy to see that if the set of states $\{|\psi_x\rangle\}$ has strong symmetries, probabilistic processes do not lead to any improvement (cf. Methods). A sketch of the argument is the following: any probabilistic process can be decomposed into the application of a filter—that transforms an input state $|\psi_x\rangle$ into the output state $|\phi_x\rangle = M_{\text{yes}}|\psi_x\rangle / \|M_{\text{yes}}|\phi_x\rangle\|$ for some suitable operator M_{yes} —

followed by a deterministic process. Now, if the set of states has strong symmetries, the optimal operator M_{yes} is forced to be equal to the identity, and the probabilistic process becomes equivalent to a deterministic one. This is the case for cloning of arbitrary pure states¹⁸, coherent states¹⁹ or coherent spin states²⁰, where the replication rates are bound to satisfy the SQL.

The HL for information replication. We now restrict our attention to the replication of states of the form $|\psi_t\rangle = e^{-itH}|\psi\rangle$, $t \in \mathbb{R}$, where $H = H^\dagger$ is a suitable Hamiltonian. We call these states clock states, as they can be generated through a time evolution obeying the Schrödinger equation. For simplicity, we focus on finite-dimensional quantum systems and we ignore the uninteresting case where H is a multiple of the identity. The main result here is a HL for the probabilistic replication of clock states. To get the result, we first establish a precision limit for probabilistic metrology where one is allowed to take advantage of filters^{21,22}. Our strategy is to apply the quantum Cramér-Rao bound (CRB)^{7,23,24} to the states $|\phi_t\rangle = M_{\text{yes}}|\psi_t\rangle / \|M_{\text{yes}}|\psi_t\rangle\|$ emerging from the filter. By explicit calculation (see Supplementary Note 2), we bound the variance at a given point t as

$$V_t^{\text{prob}} \geq \frac{1}{4(\langle \phi_t | K_t^\dagger K_t | \phi_t \rangle - \langle \phi_t | K_t | \phi_t \rangle^2)}, \quad (1)$$

where $K_t = M_{\text{yes}} H M_{\text{yes}}^{-1} + \frac{\langle \psi_t | [H, M_{\text{yes}}^\dagger M_{\text{yes}}] | \psi_t \rangle}{2\|M_{\text{yes}}|\psi_t\rangle\|^2}$. This innocent-looking application of the CRB leads immediately to a surprise: optimizing over all filters M_{yes} , the right hand side of equation (1) can be made arbitrarily small, suggesting the possibility of unlimited precision (*cf.* Methods). Note, however, that to attain the equality in equation (1), one should adapt the choice of filter to the value of t , the unknown parameter that one is trying to estimate. In practice, this is not a realistic scenario. In the case of periodic evolution, a more realistic setting is to have t distributed according to a uniform prior $p(t)$ over the period. In this case, the precision that can be achieved on average is still limited: denoting by V_p^{prob} the expected variance when the particles pass the filter, in Supplementary Note 3 we prove the bound

$$V_p^{\text{prob}} \geq \frac{1}{(E_{\text{max}} - E_{\text{min}})^2}, \quad (2)$$

where E_{max} (E_{min}) is the maximum (minimum) eigenvalue of the energy such that $|\psi\rangle$ has non-zero overlap with the corresponding eigenspace. We note that $(E_{\text{max}} - E_{\text{min}})^2/4$ is equal to the maximum of the variance of H over all possible states $|\phi\rangle$ contained in the subspace generated by the input states $\{|\psi_t\rangle\}$. Such maximum is achieved by the ‘NOON state’ $|\phi\rangle = (|E_{\text{max}}\rangle + |E_{\text{min}}\rangle)/\sqrt{2}$, where $|E_{\text{max}}\rangle$ ($|E_{\text{min}}\rangle$) is an eigenstate of H corresponding to the eigenvalues E_{max} (E_{min}). Since these states are the best states for deterministic metrology²⁵, our result implies that the best strategy for probabilistic metrology with uniform prior consists just in using a filter that generates the NOON state, and then applying the optimal deterministic estimation strategy. Up to a small correction, the same result of equation (2) can be obtained when the evolution is non-periodic, by approximating the uniform distribution with a Gaussian with large variance (*cf.* Supplementary Note 4).

When N identical copies of $|\psi_t\rangle$ are available, equation (2) gives the Heisenberg scaling

$$V_{p,N}^{\text{prob}} \geq \frac{1}{N^2(E_{\text{max}} - E_{\text{min}})^2}. \quad (3)$$

Equipped with this bound, we can now derive the HL for quantum replication; no physical process can reliably replicate a

set of clock states at a rate larger than 2. Here is a sketch of the proof: the filter transforms the product state $|\psi_t\rangle^{\otimes N}$ into the entangled state $M_{\text{yes}}|\psi_t\rangle^{\otimes N} / \|M_{\text{yes}}|\psi_t\rangle^{\otimes N}\|$. For this state, equation (3) gives the bound $V_{p,N}^{\text{prob}} \geq c/N^2$ with $c = (E_{\text{max}} - E_{\text{min}})^{-2} + \epsilon$ ($\epsilon = 0$ in the case of periodic evolution). Now, suppose that there exists a sequence of probabilistic processes with replication rate $\alpha > 1$, and suppose that the N th process in the sequence produces $M \geq N + \alpha N^\alpha$ approximate copies, with a trace distance from the ideal target vanishing asymptotically as $O(1/M^\beta)$, $\beta \geq 1$ (this assumption is lifted in Supplementary Note 5). Then, the variance of the estimation from the output state $\rho_{M,t}^{\text{out}}$ —denoted by $V_{M,t}^{\text{out}}$ —will be close to the variance of the estimation from $|\psi_t\rangle^{\otimes M}$: for every t , we will have $V_{M,t}^{\text{out}} \leq V_{M,t} + \gamma/M^\beta$, for some constant γ independent of t . Taking the average variance over the uniform prior $p(t)$, we obtain the relation $V_{M,p}^{\text{out}} \leq V_{M,p} + \gamma/M^\beta$, where $V_{M,p}^{\text{out}}$ and $V_{M,p}$ are the averages of $V_{M,t}^{\text{out}}$ and $V_{M,t}$ respectively. By definition, the average variance $V_{M,p}$ is lower bounded by the SQL and the bound is attained in the large M limit by choosing a suitable measurement (*cf.* Methods). Hence, by Taylor expanding the terms in M and keeping the leading order terms, we obtain $cN^{-2} \leq 4N^{-\alpha} + O(N^{-\alpha})$, which implies $\alpha \leq 2$. Note that here β disappeared from the equation as there is no upper bound on β , in principle the error can vanish faster than any polynomial! The HL leaves lots of room for replicating quantum information: for every rate $1 < \alpha < 2$, one has the chance not only to duplicate the input copies but also to produce an overwhelming number of replicas, with an error that vanish faster than any polynomial. In the next paragraph, we will exhibit explicit protocols that have all these features. Our protocols are necessarily probabilistic, as a consequence of the SQL, which constrains deterministic processes to produce a negligible number of extra replicas.

Probabilistic super-replication. Here we show that clock states can be reliably replicated at any rate allowed by the HL. When the rate is larger than (or equal to) 1, we call the process super-replication to emphasize the fact that it beats the SQL. The key idea to achieve super-replication is to devise a filter that modulates the Fourier amplitudes of the wavefunction in a way that enhances the replication performances to the maximum rate allowed by quantum mechanics. To move to the Fourier picture, we express the input states as $|\psi_t\rangle = \sum_{E \in \text{Spec}(H)} \sqrt{p_E} e^{-itE} |E\rangle$, where $\text{Spec}(H)$ denotes the spectrum of the Hamiltonian, $|E\rangle$ is an eigenvector of H with eigenvalue E and p_E is the probability that a measurement of the energy gives outcome E . When N identical copies are given, the joint state can be expressed as

$$|\psi_t\rangle^{\otimes N} = \sum_{E \in \text{Spec}(H^{(N)})} e^{-itE} \sqrt{p_{N,E}} |N, E\rangle, \quad (4)$$

where $H^{(N)} = \sum_{i=1}^N H_i$ is the total Hamiltonian (H_i denoting the operator H acting on the i th system), $|N, E\rangle$ is an eigenvector of $H^{(N)}$ for the eigenvalue E and $p_{N,E}$ is the probability that a measurement of the total energy gives outcome E . Choosing a filter M_{yes} that is diagonal in the energy eigenbasis, we can modulate the Fourier amplitudes of the state (4) in any way we like. Of course, since we aim at producing M perfect copies, the natural choice is to replace N with M in the probability distribution $p_{N,E}$. However, the spectrum of $H^{(N)}$ may not be contained in the spectrum of $H^{(M)}$ and in general one needs to shift the energy values by a suitable amount $\delta E_0 \approx (M - N) \langle \psi | H | \psi \rangle$ (see Supplementary Note 6). With this choice, the state after the filter is projected to the entangled state $|\Phi_t^N\rangle \propto \sum_{E \in \text{Spec}(H^{(N)})} e^{-itE} \sqrt{p_{M,E + \delta E_0}} |N, E\rangle$.

The state $|\Phi_t^N\rangle$ will now act as a quantum programme, containing the instructions that will be used by a deterministic quantum device to generate M approximate clones. For this purpose, we use a device that coherently transforms each eigenstate $|N, E\rangle$ into the corresponding eigenstate $|M, E + \delta E_0\rangle$, thus producing the state $|\Psi_t^M\rangle \propto \sum_{E \in \text{Spec}(H^{(N)})} e^{-iEt} \sqrt{P_{M, E + \delta E_0}} |M, E + \delta E_0\rangle$. It is not hard to see that the state $|\Psi_t^M\rangle$ has high fidelity with the desired state $|\psi_t\rangle^{\otimes M}$ for every replication rate allowed by the HL. Precisely, for large enough N one has the bound

$$F_{N \rightarrow M} \geq 1 - 2K \exp\left(-\frac{2p_{\min}^2 N^2}{M} + \frac{4N}{MK}\right), \quad (5)$$

where K is the number of energy levels E such that $p_E \neq 0$ and $p_{\min} = \min\{p_E : p_E \neq 0\}$ (see Supplementary Note 7 for the proof). Equation (5) shows that the fidelity approaches 1 faster than any polynomial whenever M is of order N^2 with $\alpha < 2$, that is, whenever the replication rate satisfies the HL. A strong converse can be proven: every process replicating quantum states at a rate higher than 2 must have vanishing fidelity in the limit of large N , as showed in Supplementary Note 5.

Thanks to our filter, we have been able to embezzle from nature a large number of replicas. The improvement is striking if one compares it with the performances of standard cloning processes. Let us illustrate this fact for the replication of linear polarization states $|\psi_t\rangle = \cos t|V\rangle + \sin t|H\rangle$. In this case, the best deterministic process is the phase-covariant cloner of ref. 26, and its fidelity is $F_{N \rightarrow M}^{\text{det}} \approx \frac{2\sqrt{MN}}{M+N}$ in the asymptotic limit of large N and M . In agreement with our SQL, the fidelity vanishes whenever M is of order $N^{1+\epsilon}$, $\epsilon > 0$. As we proved in general, deterministic processes can only produce a number of extra replicas that is a negligible fraction of the number of input photons. In stark contrast, for $N = 100$ input photons, our filter can produce $M = 1,000$ approximate copies with fidelity

$F_{\text{prob}} = 0.9986$, whereas the fidelity of the optimal deterministic cloner is only $F_{\text{det}} = 0.5739$. In the example of linearly polarized photons, our filter is provenly optimal, as it coincides with the optimal probabilistic cloner of ref. 17. Quite surprisingly, the possibility of super-replication was not recognized in ref. 17 where the advantage of probabilistic processing was conjectured to be only of the order of one percent.

The different features of replication processes at the HL and at the SQL are illustrated in Fig. 1 in the case of linearly polarized photons.

The advantage of the filter is clear also in the non-asymptotic setting: a plot comparing the performances of replication with and without the filter in the case $N = 20$ is presented in Fig. 2.

What makes the improvement even more dramatic is that the quality of the replicas is measured by the global fidelity between the output state and the desired joint state of M perfect copies, which is much smaller than the fidelity that each single copy has with the state $|\psi_t\rangle$.

Maximizing the probability of success. The dazzling performances of super-replication come at a price: the probability that the input systems pass the filter has to decay with N . Indeed, in Supplementary Note 1 we already showed a strong converse of the SQL: every deterministic process producing replicas at a rate higher than the SQL must have vanishing fidelity.

For our filter, the probability of super-replication at rate $\alpha > 1$ decreases exponentially fast as $p_{\text{yes}}[N \rightarrow M] \leq e^{-kN}$ for a suitable constant $k > 0$ independent of α (see Supplementary Note 8). One may ask whether there are ways to achieve super-replication with a larger success probability. It turns out that the answer is yes. In Supplementary Note 9, we show that a process producing $M \leq O(N)$ replicas can have a success probability $p_{\text{yes}}[N \rightarrow M]$ going to zero as $1/N^\delta$ for every desired $\delta > 0$. Most importantly, for a super-replication process with rate $\alpha = 1 + \epsilon$, $\epsilon > 0$, we show




Figure 1 | The advantages of probabilistic super-replication. (a) Every deterministic process is constrained by the SQL: the number of clones that can be produced reliably from N input copies is of order N and the fidelity of the clones cannot approach 1 faster than the inverse of a polynomial of degree 4. (b) Super-replication at the HL: the cloning performances can be dramatically increased by a probabilistic filter, depicted here as a ‘quantum Maxwell demon’ that separates two branches of the wavefunction. In the successful branch, any number of clones of order $N^{2-\epsilon}$ can be produced with fidelity approaching 1 faster than any polynomial.




Figure 2 | Replication of linearly polarized photons. The fidelity of the best cloning processes for linearly polarized photons $|\psi_t\rangle = \cos t|V\rangle + \sin t|H\rangle$ is plotted here for $N=20$ input copies, with the number M of output clones ranging from 20 to 400. The red (green) line refers to optimal cloning with (without) filter.

how to increase the success probability to $p_{\text{yes}}[N \rightarrow M] \geq e^{-N\delta}$ for every desired exponent $\delta > \epsilon$. However, no further improvement is possible below the critical value $\delta_c = \epsilon$; any process with success probability scaling as $p_{\text{yes}}[N \rightarrow M] = ae^{-bN\delta}$, $\delta < \epsilon$, $a, b > 0$ must have vanishing fidelity in the asymptotic limit, as observed in Supplementary Note 1.

These results identify the optimal exponent for the decay of the success probability in a super-replication process, pinning down the trade-off between the replication rate and success probability. Interestingly, the choice between the advantage of a high replication rate and that of a unit success probability can be always delayed to the very last moment. Indeed, it is easy to see that asymptotically the best deterministic replication process is just given by the coherent transformation $|N, E\rangle \mapsto |M, E + \delta E_0\rangle$ that we used in our protocol. Hence, an alternative way to achieve super-replication is to apply first the best deterministic process and later to modulate the Fourier amplitudes of the wavefunction using a probabilistic filter.

Many-worlds fairness. Super-replication can be achieved not only for one-parameter families of clock states but also for different manifolds of states, including the manifold of all maximally entangled states of two identical systems. However, super-replication is not a generic feature. As we already mentioned, if one tries to copy an arbitrary—as opposed to linear—polarization state, then no filter is going to help; the performances with filter are equal to the performances without filter. More generally, for generic quantum systems, no probabilistic filter can make an arbitrary quantum state more copiable. This property is quite compelling when considered from the angle of the many-worlds interpretation of quantum mechanics²⁷ because it states that no-branch of the wavefunction of the universe offers an advantage over the others in replicating the information contained in a completely unknown state. Regarding different branches as ‘different worlds’, we can formulate this as a fundamental principle, which we name many-worlds fairness: the maximum rate at which arbitrary information can proliferate is the same in all possible worlds. Many world fairness rules out quantum

mechanics on real Hilbert spaces²⁸, an alternative physical theory where photons can have only linear polarizations. Thanks to this observation, one can provide a new answer to the old question ‘why are physical systems described by complex—instead of real—Hilbert spaces?’. Traditionally, the standard answer has been to invoke local tomography, the property that one can completely identify a mixed state from the statistics of local measurement on the components. However, one may consider this as an *ad hoc* requirement²⁹, and, in fact, there are even reasons to prefer real quantum mechanics to its complex version, as it was recently pointed out by Wootters³⁰. Balancing this fact, many-world fairness offers a new reason (other than the usual local tomography) in favour of complex quantum mechanics.

Discussion

Super-replication has been introduced here from a theoretical point of view. But is it possible to implement it experimentally? Luckily, very recently there have been experimental breakthroughs on the closely related topic of probabilistic amplification of coherent states of light^{31–35}. Although the translation to our case is not immediate, we suggest that super-replication of linearly polarized photons could be achieved through a suitable sequence of amplitude-damping channels, which for the polarization play the role of the photon subtraction for coherent states. An alternative approach is to first encode the state of the input copies into a coherent state via matter-light teleportation³⁶, amplify the coherent state in a probabilistic fashion, and then teleport back. This scheme provides a new application of the existing experimental schemes for coherent state amplification, making them the building block for the replication of quantum information at the HL. Finally, a third avenue towards implementing super-replication would be through stimulated emission, combined with a suitable monitoring of the emitting atoms. Stimulated emission has been discussed extensively in connection with deterministic cloning^{37–39}, but its potential for implementing probabilistic processes is still unexplored.

In relation to the existing literature, it is interesting to comment on the relation between our results and previous works on probabilistic cloning and probabilistic estimation. The idea that in some situations even non-orthogonal states can be copied perfectly using a probabilistic device was first introduced by Duan and Guo⁴⁰ who showed that a set of pure quantum states can be copied perfectly if and only if they are linearly independent. This means that for a single photon, one can perfectly copy two polarization states, while for N photons one can copy at most $N+1$ states. Subsequent works^{17,41,42} showed that probabilistic devices can improve the performances of approximate cloning for continuous sets of states, leading to nearly ideal performances in the case of coherent states with fixed amplitude. The nearly perfect cloning of coherent states may lead one to believe that in general there is no limit to the amount of clones that can be produced probabilistically. Contrarily to this intuition, we have proven here that for finite quantum systems, the HL sets the ultimate bound $M = O(N^2)$ to the number of clones that can be produced reliably. Moreover, we have shown that for arbitrary one-parameter families of states, any replication rate allowed by the HL can be achieved and we identified exactly the exponent at which the probability of success has to decay.

Finally, the relation between cloning and estimation has been extensively investigated in the literature^{43–47} in terms of single-copy fidelities. However, none of the approaches proposed so far was suitable to derive limits on the asymptotic replication rates, nor to connect the latter with quantum metrology. Recently, it has been observed that the precision of quantum estimation can

be improved using probabilistic strategies^{21,22} where depending on the outcome of the filter, one can decide to abstain from estimating the parameter. The mechanism of abstention can sometimes boost the precision from the SQL $1/N$ to the HL $1/N^2$. The connection between these results and super-replication is made clear by our approach: the fact that a probabilistic filter can improve estimation until the HL implies that N^2 , rather than N , is the upper bound for the replication rate of the states $|\psi_t\rangle = e^{-itH}|\psi_i\rangle$. However, there is an important catch: the filter that achieves replication at the HL is not the same filter that achieves estimation at the HL. In fact, if we were to use the same filter needed in refs 21,22, we would not be able to reduce the error down to zero.

Methods

Decomposition of quantum instruments. To describe the most general probabilistic processes allowed by quantum mechanics, we use the framework of quantum instruments^{48–50}. A quantum instrument with input (output) Hilbert space \mathcal{H}_{in} (\mathcal{H}_{out}) is a collection of completely positive (CP), trace non-increasing linear maps $\{\mathcal{P}_j\}_{j \in Y}$, where each map transforms density matrices on \mathcal{H}_{in} into (sub-normalized) density matrices on \mathcal{H}_{out} . If the input system is prepared in the density matrix ρ , the probability that outcome j is $p(j|\rho) = \text{Tr}[\mathcal{P}_j(\rho)]$. Conditionally, to the occurrence of outcome j , the state of the output system is $\rho_j^{\text{out}} = \mathcal{P}_j(\rho)/\text{Tr}[\mathcal{P}_j(\rho)]$. We note an elementary fact about quantum instruments: any quantum instrument can be decomposed into a pure measurement on the input system followed by quantum channel depending on the outcome. By ‘pure measurement on the input system’, we mean a quantum instrument $\{\mathcal{M}_j\}_{j \in Y}$ of the special form $\mathcal{M}_j(\rho) = M_j \rho M_j^\dagger$ for some operator M_j on \mathcal{H}_{in} . By ‘quantum channel’, we mean a trace-preserving CP map \mathcal{C}_j , transforming density matrices on \mathcal{H}_{in} into (normalized) density matrices on \mathcal{H}_{out} . In formula, our claim is that every CP map \mathcal{P}_j can be decomposed as

$$\mathcal{P}_j = \mathcal{C}_j \mathcal{M}_j. \tag{6}$$

The proof is as follows: let \mathcal{P}_j^\dagger be the adjoint of the map \mathcal{P}_j , defined by

$$\text{Tr}[\mathcal{P}_j^\dagger(A)\rho] = \text{Tr}[A\mathcal{P}_j(\rho)] \text{ for every operator } A \text{ and for every density matrix } \rho.$$

Define the operator $M_j := \sqrt{\mathcal{P}_j^\dagger(I)}$ and the map \mathcal{L}_j by the relation $\mathcal{L}_j(\rho) := \mathcal{P}_j(M_j^{-1}\rho M_j^{-1})$, where M_j^{-1} is the inverse of M_j on its support. If M_j is invertible, then $\mathcal{C}_j = \mathcal{L}_j$ is trace preserving:

$$\begin{aligned} \text{Tr}[\mathcal{L}_j(\rho)] &= \text{Tr}[\mathcal{P}_j(M_j^{-1}\rho M_j^{-1})] \\ &= \text{Tr}\left\{\mathcal{P}_j^\dagger(I)\left[\mathcal{P}_j^\dagger(I)\right]^{-1/2}\rho\left[\mathcal{P}_j^\dagger(I)\right]^{-1/2}\right\} = \text{Tr}[\rho]. \end{aligned}$$

In this case, by definition we have $\mathcal{P}_j(\rho) = \mathcal{C}_j(M_j\rho M_j^\dagger)$ for every ρ . If M_j is not invertible, we can define the trace-preserving map $\mathcal{C}_j(\rho) := \mathcal{L}_j(\rho) + P_\perp \rho P_\perp$, where P_\perp is the projector on the kernel of M_j . Again, the definition implies $\mathcal{P}_j(\rho) =$

$$\mathcal{C}_j(M_j\rho M_j^\dagger) \text{ for every state } \rho.$$

For the purposes of this paper, it is enough to consider measurements with two outcomes (yes,no), referred to as filters. The filter induces a bifurcation of the wavefunction and selects one particular branch, corresponding to the successful outcome yes.

Symmetry constraints on probabilistic cloning. Consider a quantum cloning problem where the set of states $\{|\psi_x\rangle\}$ has a group of symmetries, denoted by G . This means that, for every unitary operator U_g representing the action of a symmetry transformation in the group, one has $\{U_g|\psi_x\rangle\}_{x \in X} = \{|\psi_x\rangle\}_{x \in X}$. A probabilistic cloner is described by a quantum instrument $\{\mathcal{P}_{\text{yes}}, \mathcal{P}_{\text{no}}\}$ from the space of N copies to that of M copies, where the CP maps \mathcal{P}_{yes} and \mathcal{P}_{no} correspond to the successful and unsuccessful instances, respectively. Conditional to the successful outcome, the fidelity of the output state with the desired M -copy state is

$$F_x^{\text{prob}}[\mathcal{P}_{\text{yes}}] = \frac{\langle \psi_x |^{\otimes M} \mathcal{P}_{\text{yes}}[(|\psi_x\rangle\langle\psi_x|)^{\otimes N}] | \psi_x \rangle^{\otimes M}}{\text{Tr}\{\mathcal{P}_{\text{yes}}[(|\psi_x\rangle\langle\psi_x|)^{\otimes N}]\}}. \tag{7}$$

The goal of cloning is to maximize the worst-case fidelity $F_{\text{wc}}^{\text{prob}}[\mathcal{P}_{\text{yes}}] := \inf_x F_x^{\text{prob}}[\mathcal{P}_{\text{yes}}]$. Because of the symmetry of the set of states, the maximization can be restricted without loss of generality to the set of covariant CP maps, that is, of CP maps satisfying the relation

$$U_g^{\otimes M}[\mathcal{P}_{\text{yes}}(\rho)]U_g^{\dagger \otimes M} = \mathcal{P}_{\text{yes}}(U_g^{\otimes N}\rho U_g^{\dagger \otimes N}), \tag{8}$$

for every group element g and for every quantum state ρ . The proof is standard and we refer the interested reader to similar proofs provided in the literature, such as

those of refs 51,52 for covariant measurements and that of ref. 53 for covariant cloning channels.

Let us work out the implication of symmetry, starting from the decomposition $\mathcal{P}_{\text{yes}} = \mathcal{C}_{\text{yes}}\mathcal{M}_{\text{yes}}$ of equation (6). It is easy to see that the covariance of \mathcal{P}_{yes} implies the commutation relation $[\mathcal{M}_{\text{yes}}, U_g^{\otimes N}] = 0$, for every group element g . Indeed, using the definition $M_{\text{yes}} := \sqrt{\mathcal{P}_{\text{yes}}^\dagger(I)}$ and the covariance of $\mathcal{P}_{\text{yes}}^\dagger$, we have

$$\begin{aligned} U_g^{\otimes N}M_{\text{yes}}U_g^{\dagger \otimes N} &= U_g^{\otimes N}\sqrt{\mathcal{P}_{\text{yes}}^\dagger(I)}U_g^{\dagger \otimes N} \\ &= \sqrt{\mathcal{P}_{\text{yes}}^\dagger(U_g^{\otimes M}U_g^{\dagger \otimes M})} \\ &= \sqrt{\mathcal{P}_{\text{yes}}^\dagger(I)} \\ &= M_{\text{yes}}. \end{aligned}$$

Now, if the group of symmetries is sufficiently large, the action of the unitaries $U_g^{\otimes N}$ can be irreducible in the subspace containing the input states $|\psi_x\rangle^{\otimes N}$. When this happens, the Schur’s lemma imposes that the filter M_{yes} be a multiple of the identity, and, therefore $\mathcal{P}_{\text{yes}} \propto \mathcal{C}_{\text{yes}}$. By equation (7), this means that the fidelity achieved by the CP map \mathcal{P}_{yes} is the same as the fidelity achieved by the quantum channel \mathcal{C}_{yes} . In summary, we have proven the following fact: the maximum of the fidelity over all probabilistic cloners is equal to the maximum of the fidelity over deterministic cloners whenever the action of the unitaries $\{U_g^{\otimes N}\}$ is irreducible in the subspace spanned by the input states $\{|\psi_x\rangle^{\otimes N}\}$. This simple observation has several important consequences. A first consequence, noted by Fiuřářek¹⁷, is that the probabilistic processes do not offer any advantage for 1-to- M cloning of arbitrary pure states. Our approach allows to reach the same conclusion in a fairly broader range of cloning problems: in particular, it implies that probabilistic processes offer no advantage in the case of N -to- M cloning of arbitrary pure states¹⁸, N -to- M cloning of coherent states¹⁹ and spin-coherent states²⁰, 1-to- M cloning of two Fourier transformed bases⁵³ and of all the sets of states considered in ref. 53, and 1-to- M phase-covariant cloning for qubit states on the equator of the Bloch sphere²⁶.

Pointwise versus average CRB. The CRB is the cornerstone of quantum metrology^{5–9,11}. For a family of clock states $|\psi_t\rangle = e^{-itH}|\psi\rangle$, it states that the variance in the estimation of t from the state $|\psi_t\rangle$ is lower bounded by $V_t \geq 1/Q_t$, where $Q_t = 4(\langle\psi_t|H^2|\psi_t\rangle - \langle\psi_t|H|\psi_t\rangle^2)$ is the quantum Fisher information (QFI). The equality in the CRB can be achieved if one adapts the estimation strategy to the value of t . This can be done in an asymptotic scenario when one is given a large number of copies of $|\psi_t\rangle$ and uses a fraction of them to obtain a rough estimate of t (ref. 54). However, if instead of having N copies one has a single N -particle entangled state, the CRB may not give an achievable lower bound. For example, Hayashi⁵⁵ showed that the CRB for phase estimation predicts the correct scaling c/N^2 , but with a constant c that is not achievable. In the case of probabilistic metrology, the issue about the achievability of the CRB appears in a more dramatic way: if we allow one to adapt also the choice of the filter M_{yes} to the value of t , then the CRB promises unlimited precision. It is instructive to see how this phenomenon can happen. For a given t_0 , choose a parametrization where $\langle\psi_{t_0}|H|\psi_{t_0}\rangle = 0$ (this can always be done by redefining H as $H' = H - \langle\psi_{t_0}|H|\psi_{t_0}\rangle$). Then, define the vector

$$|\psi_{t_0}^\perp\rangle := \frac{H|\psi_{t_0}\rangle}{\|H|\psi_{t_0}\rangle\|},$$

which by construction satisfies $\langle\psi_{t_0}^\perp|\psi_{t_0}\rangle = 0$. Now, choose the filter operator

$$M_{\text{yes},t_0} := \epsilon|\psi_{t_0}\rangle\langle\psi_{t_0}| + |\psi_{t_0}^\perp\rangle\langle\psi_{t_0}^\perp|$$

The filter does not change the state $|\psi_{t_0}\rangle$:

$$|\phi_{t_0}\rangle = \frac{M_{\text{yes},t_0}|\psi_{t_0}\rangle}{\|M_{\text{yes},t_0}|\psi_{t_0}\rangle\|} \equiv |\psi_{t_0}\rangle.$$

However, it amplifies the time derivative at t_0 :

$$\begin{aligned} \left.\frac{d|\phi_t\rangle}{dt}\right|_{t=t_0} &= K_{t_0}|\phi_{t_0}\rangle \\ &= \left(M_{\text{yes},t_0}HM_{\text{yes},t_0}^{-1}\right)|\psi_{t_0}\rangle \\ &= \frac{H|\psi_{t_0}\rangle}{\epsilon} \\ &= \frac{1}{\epsilon}\left.\frac{d|\psi_t\rangle}{dt}\right|_{t=t_0}, \end{aligned}$$

where in the second line we used the explicit expression

$$K_{t_0} = M_{\text{yes},t_0}HM_{\text{yes},t_0}^{-1} + \frac{\langle\psi_{t_0}|[H, M_{\text{yes},t_0}^\dagger M_{\text{yes},t_0}]\psi_{t_0}\rangle}{2\|M_{\text{yes}}|\psi_{t_0}\rangle\|^2},$$

along with the observation that $\langle \psi_{t_0} | [H, M_{\text{yes}, t_0}^\dagger] | \psi_{t_0} \rangle = 0$. As a result, the probabilistic QFI at t_0 , defined as the QFI associated to the states $\{|\phi_t\rangle\}$, satisfies

$$Q_{t_0}^{\text{prob}} = \langle \phi_{t_0} | K_{t_0}^\dagger K_{t_0} | \phi_{t_0} \rangle - \langle \phi_{t_0} | K_{t_0} | \phi_{t_0} \rangle^2 \\ = \frac{Q_{t_0}}{\epsilon^2}$$

This equation means that, whenever the original QFI is non-zero at a specific value t_0 , the probabilistic QFI can be made arbitrarily large with a suitable choice of ϵ . This fact, combined with the CRB, seems to suggest that one can have unlimited precision. However, such a conclusion is an artifact of the pointwise character of the CRB: the probabilistic estimation scheme shown above has unlimited precision only in an infinitesimally small neighbourhood of t_0 . In light of this observation, for probabilistic metrology, it is more sensible to choose the filter M_{yes} independent of t , to assign a prior $p(t)$ to the unknown parameter and to minimize the expected variance. Note that, since the estimate is produced only when the system passes through the filter, the expectation has to be computed with respect to the conditional probability distribution $p(t | \text{yes}) \propto \|M_{\text{yes}}|\psi_t\rangle\|^2 p(t)$. The expected variance, denoted by $V_p^{\text{prob}} := \int dt p(t | \text{yes}) V_t$, is lower bounded by the inverse of the expected QFI, denoted by $Q_p^{\text{prob}} := \int dt p(t | \text{yes}) Q_t$ in formula

$$V_p^{\text{prob}} \geq \frac{1}{Q_p^{\text{prob}}}. \quad (9)$$

The proof of equation (9) is an elementary application of the pointwise quantum CRB and of the Schwartz inequality:

$$1 \leq \left[\int dt p(t | \text{yes}) \sqrt{V_t Q_t} \right]^2 \\ \leq \left[\int dt p(t | \text{yes}) V_t \right] \left[\int dt p(t | \text{yes}) Q_t \right] \\ \equiv V_p^{\text{prob}} Q_p^{\text{prob}}.$$

Note that the average CRB of equation (9) is just a lower bound and the equality may not be achievable. However, achievability is of no consequence for our arguments about the replication rates.

References

- Wootters, W. & Zurek, W. A single quantum cannot be cloned. *Nature* **299**, 802–803 (1982).
- Dieks, D. Communication by EPR devices. *Phys. Lett. A* **92**, 271–272 (1982).
- Herbert, N. FLASH—A superluminal communicator based upon a new kind of quantum measurement. *Found. Phys.* **12**, 1171–1179 (1982).
- Einstein, A. Strahlungs-emission und -absorption nach der Quantentheorie. *Verh. Dtsch Phys. Ges.* **18**, 318–323 (1916).
- Wineland, D. J., Bollinger, J. J., Itano, W. M. & Moore, F. L. Spin squeezing and reduced quantum noise in spectroscopy. *Phys. Rev. A* **46**, R6797–R6800 (1992).
- Holland, M. J. & Burnett, K. Interferometric detection of optical phase shift at the Heisenberg limit. *Phys. Rev. Lett.* **71**, 1355–1358 (1993).
- Braunstein, S. L. & Caves, C. M. Statistical distance and the geometry of quantum states. *Phys. Rev. Lett.* **72**, 3439–3443 (1994).
- Giovannetti, V., Lloyd, S. & Maccone, L. Quantum-enhanced measurements: beating the standard quantum limit. *Science* **306**, 1330–1336 (2004).
- Giovannetti, V., Lloyd, S. & Maccone, L. Advances in quantum metrology. *Nat. Photonics* **5**, 222–229 (2011).
- Huelga, S. F. *et al.* Improvement of frequency standards with quantum entanglement. *Phys. Rev. Lett.* **79**, 3865–3868 (1997).
- Escher, B. M., de Matos Filho, R. L. & Davidovich, L. General framework for estimating the ultimate precision limit in noisy quantum-enhanced metrology. *Nat. Phys.* **7**, 406–411 (2011).
- Bužek, V. & Hillery, M. Quantum copying: beyond the no-cloning theorem. *Phys. Rev. A* **54**, 1844–1852 (1996).
- Scarani, V., Iblisdir, S., Gisin, N. & Acin, A. Quantum cloning. *Rev. Mod. Phys.* **77**, 1225 (2005).
- Cerf, N. J. & Fiurášek, J. Optical quantum cloning. *Prog. Optics* **49**, 455–545 (2006).
- Braginsky, V. B. & Khalili, F. Y. *Quantum Measurement* (Cambridge University Press, 1995).
- Nielsen, M. A. & Chuang, I. L. *Quantum Computation and Quantum Information* (Cambridge University Press, 2000).
- Fiurášek, J. Optimal probabilistic cloning and purification of quantum states. *Phys. Rev. A* **70**, 032308 (2004).
- Werner, R. F. Optimal cloning of pure states. *Phys. Rev. A* **58**, 1827–1832 (1998).
- Cerf, N. J., Ipe, A. & Rottenberg, X. Cloning of continuous quantum variables. *Phys. Rev. Lett.* **85**, 1754–1757 (2000).
- Demkowicz-Dobrzański, R., Kus, M. & Wodkiewicz, K. Cloning of spin-coherent states. *Phys. Rev. A* **69**, 012301 (2004).
- Fiurášek, J. Optimal probabilistic estimation of quantum states. *N. J. Phys.* **8**, 192 (2006).
- Gendra, B., Ronco-Bonvehi, E., Calsamiglia, J., Muñoz-Tapia, R. & Bagan, E. Quantum metrology assisted with abstention. *Phys. Rev. Lett.* **110**, 100501 (2013).

- Helstrom, C. W. *Quantum Detection and Estimation Theory* (Academic Press, 1976).
- Fujiwara, A. & Nagaoka, H. Quantum Fisher metric and estimation for pure state models. *Phys. Lett. A* **201**, 119–124 (1995).
- Giovannetti, V., Lloyd, S. & Maccone, L. Quantum metrology. *Phys. Rev. Lett.* **96**, 010401 (2006).
- D’Ariano, G. & Macchiavello, C. Optimal phase-covariant cloning for qubits and qutrits. *Phys. Rev. A* **67**, 042306 (2003).
- Everett, H. Relative state formulation of quantum mechanics. *Rev. Mod. Phys.* **29**, 454–462 (1957).
- Stueckelberg, E. C. G. Quantum theory in real Hilbert space. *Helv. Phys. Acta* **33**, 458 (1960).
- Hardy, L. & Wootters, W. Limited holism and real-vector-space quantum theory. *Found. Phys.* **3**, 454–473 (2012).
- Wootters, W. *Optimal Information Transfer and Real-Vector-Space Quantum Theory*. Preprint at <http://arXiv.org/abs/1301.2018> (2013).
- Xiang, G. Y., Ralph, T. C., Lund, A. P., Walk, N. & Pryde, G. J. Heralded noiseless linear amplification and distillation of entanglement. *Nat. Photonics* **4**, 316–319 (2010).
- Ferreyl, F. *et al.* Implementation of a nondeterministic optical noiseless amplifier. *Phys. Rev. Lett.* **104**, 123603 (2010).
- Usuga, M. A. *et al.* Noise-powered probabilistic concentration of phase information. *Nat. Phys.* **6**, 767–771 (2010).
- Zavatta, A., Fiurášek, J. & Bellini, M. A high-fidelity noiseless amplifier for quantum light states. *Nat. Photonics* **5**, 52–60 (2011).
- Kocsis, S., Xiang, G. Y., Ralph, T. C. & Pryde, G. F. Heralded noiseless amplification of a photon polarization qubit. *Nat. Phys.* **9**, 23–28 (2013).
- Sherson, J. F. *et al.* Quantum teleportation between light and matter. *Nature* **443**, 557–560 (2006).
- Milonni, P. W. & Hardies, M. L. Photons cannot always be replicated. *Phys. Lett. A* **92**, 321–322 (1982).
- Mandel, L. Is a photon amplifier always polarization dependent? *Nature* **304**, 188 (1983).
- Simon, C., Weihs, G. & Zeilinger, A. Optimal quantum cloning via stimulated emission. *Phys. Rev. Lett.* **84**, 2993–2996 (2000).
- Duan, L. M. & Guo, G. C. A probabilistic cloning machine for replicating two non-orthogonal states. *Phys. Rev. Lett.* **80**, 4999–5002 (1998).
- Ralph, T. C. & Lund, A. P. Nondeterministic noiseless linear amplification of quantum systems. *Quantum Communication Measurement and Computing Proceedings of 9th International Conference*. (ed Lvovsky, A.) 155–160 (AIP, 2009).
- Chiribella, G. & Xie, J. Optimal design and quantum benchmarks for coherent state amplifiers. *Phys. Rev. Lett.* **110**, 213602 (2013).
- Gisin, N. & Massar, S. Optimal quantum cloning machines. *Phys. Rev. Lett.* **79**, 2153–2156 (1997).
- Bruss, D., Ekert, A. & Macchiavello, C. Optimal universal quantum cloning and state estimation. *Phys. Rev. Lett.* **81**, 2598–2601 (1998).
- Bae, J. & Acin, A. Asymptotic quantum cloning is state estimation. *Phys. Rev. Lett.* **97**, 030402 (2006).
- Chiribella, G. & D’Ariano, G. Quantum information becomes classical when distributed to many users. *Phys. Rev. Lett.* **97**, 250503 (2006).
- Chiribella, G. On quantum estimation, quantum cloning, and finite quantum de Finetti theorems. *Theory of Quantum Computation, Communication, and Cryptography, Lecture Notes in Computer Science* (Springer, 2011).
- Davies, E. B. & Lewis, J. T. An operational approach to quantum probability. *Comm. Math. Phys.* **17**, 239–260 (1970).
- Ozawa, M. Quantum measuring processes of continuous observables. *J. Math. Phys.* **25**, 79 (1984).
- Busch, P., Grabowski, M. & Lahti, P. J. *Operational Quantum Physics* Vol. 31 (Springer, 1995).
- Holevo, A. S. *Probabilistic and Statistical Aspects of Quantum Theory* (Springer 1982).
- Ozawa, M. in *Research Reports on Information Sciences, Series A: Mathematical Sciences* 74 (Department of Information Sciences, Tokyo Institute of Technology, 1980).
- Chiribella, G., D’Ariano, G. M., Perinotti, P. & Cerf, N. J. Extremal quantum cloning machines. *Phys. Rev. A* **72**, 042336 (2005).
- Gill, R. D. & Massar, S. State estimation for large ensembles. *Phys. Rev. A* **61**, 042312 (2000).
- Hayashi, M. Comparison between the Cramér-Rao and the mini-max approaches in quantum channel estimation. *Comm. Math. Phys.* **304**, 689–709 (2011).

Acknowledgements

This work is supported by the National Basic Research Program of China (973) 2011CBA00300 (2011CBA00301) and by the National Natural Science Foundation of China through Grants 11350110207, 61033001 and 61061130540, and by the 1000 Youth Fellowship Program of China. We are grateful to the two anonymous Referees for comments that stimulated a substantial strengthening of the original manuscript. A particular acknowledgement goes to Xinhui Yang for drawing Fig. 1.

Author contributions

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

Additional information

Supplementary information accompanies this paper at <http://www.nature.com/naturecommunications>.

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

Reprints and permission information is available online at <http://npg.nature.com/reprintsandpermissions/>

How to cite this article: Chiribella, G. *et al.* Quantum replication at the Heisenberg limit. *Nat. Commun.* 4:2915 doi: 10.1038/ncomms3915 (2013).