Editors' Suggestion Featured in Physics

No-Go Theorems for Quantum Resource Purification

Kun Fang[1,2,*] and Zi-Wen Liu[3,†]

[1]Department of Applied Mathematics and Theoretical Physics, University of Cambridge, Cambridge, CB3 0WA, United Kingdom
[2]Institute for Quantum Computing, University of Waterloo, Waterloo, Ontario N2L 3G1, Canada
[3]Perimeter Institute for Theoretical Physics, Waterloo, Ontario N2L 2Y5, Canada

(Received 17 September 2019; revised 26 February 2020; accepted 30 June 2020; published 6 August 2020)

The manipulation of quantum "resources" such as entanglement, coherence, and magic states lies at the heart of quantum science and technology, empowering potential advantages over classical methods. In practice, a particularly important kind of manipulation is to "purify" the quantum resources since they are inevitably contaminated by noise and thus often lose their power or become unreliable for direct usage. Here we prove fundamental limitations on how effectively generic noisy resources can be purified enforced by the laws of quantum mechanics, which universally apply to any reasonable kind of quantum resource. More explicitly, we derive nontrivial lower bounds on the error of converting any full-rank noisy state to any target pure resource state by any free protocol (including probabilistic ones)—it is impossible to achieve perfect resource purification, even probabilistically. Our theorems indicate strong limits on the efficiency of distillation, a widely used type of resource purification routine that underpins many key applications of quantum information science. In particular, this general result induces the first explicit lower bounds on the resource cost of magic state distillation, a leading scheme for realizing scalable fault-tolerant quantum computation. Implications for the standard error-correction-based methods are specifically discussed.

DOI: 10.1103/PhysRevLett.125.060405

The field of quantum information takes a pragmatic approach to examining and utilizing quantum mechanics, seeking to obtain rigorous understandings of which information processing tasks can or cannot be accomplished according to the laws of nature. Efforts along this line since the 1980s have revolutionized our perception of physics and paved the way for many innovative technological applications such as quantum computation and communication [1,2]. In particular, the formulations of no-go (impossibility) theorems have played seminal roles—they often represent key advances in our understanding of quantum mechanics and have exerted profound influence on the development of quantum information science and technology. A representative example is the no-cloning theorem [3,4], which directly led to the invention of quantum error correction [5,6] and laid the foundation for plenty of other major quantum applications such as quantum cryptography [7], as well as advancing our understanding of the foundations of quantum mechanics [8,9].

At the heart of the desired quantum information processing tasks is the manipulation of various useful quantum features, the most prominent examples being entanglement [10], coherence [11], and "magic" [12,13], that emerge as valuable "resources" that are needed to empower advantages over classical methods. Such resource features can arise from all kinds of physical or conceptual restrictions on the feasible operations. A prototypical example is the "distant labs" paradigm where only local operations within the separate labs and classical communication between them (the so-called "LOCC") is allowed, rendering entanglement a resource that cannot be obtained for free and could, for instance, enable efficient quantum communication [2,10].

In practice, a particularly important and widely studied kind of manipulation is to "purify" the quantum resources, since quantum systems are highly susceptible to faulty controls and noise effects such as decoherence [2,14] that may jeopardize the power and reliability of quantum resources. In particular, a standard procedure of quantum resource purification is to extract high-quality resource states better suited for application from a large amount of raw noisy ones, which is known as *distillation*. Most notably, the distillation of entanglement [15–17], coherence [18–20], and magic states [12] has been extensively studied as a key subroutine in quantum computation and communication. Therefore, understanding the limits to the efficiency of purification and distillation tasks is of great theoretical and practical importance.

To address this problem in a rigorous and general manner, we shall use the language of quantum resource theory (see [21] for an introduction of this framework), where each resource theory is defined by a set of *free states* (in contrast to *resource states*) and a set of *free operations*. Again, take the entanglement theory as an example: the set of free states consists of the separable (unentangled) states, and LOCC is a standard choice of the set of free operations. Free states and operations can be flexibly defined, which gives rise to a wide variety of meaningful resource theories as long as they follow a *golden rule*: any free operation can only map a free state to another free state. This simple rule selects the largest possible set of free operations allowed in resource manipulation since any other operation can by definition create resources and thus trivialize the theory. Moreover, note that we are interested in the *one-shot* setting, as opposed to the conventional asymptotic setting here, since only a finite amount of resources is accessible in reality. We refer readers to Ref. [22] for a general theory of the rates of one-shot resource manipulation.

In this work, we prove a set of no-go theorems for quantum resource purification that universally apply to any reasonable resource theory, manifesting that the production of any pure resource state with an arbitrarily small error, however weak this target state is, is generically prohibited by the golden rule. More formally, we establish quantitative bounds on the achievable accuracy of any free operation that is supposed to work with some probability. It turns out that there is a nontrivial trade-off between the accuracy and success probability akin to the uncertainty relations. The proofs follow from analyzing the peculiar properties of the hypothesis testing relative entropy monotone, a quantity known to characterize the efficiency of one-shot distillation in many cases [20,22–27] but not studied in great depth. Using the above results, we find lower bounds on the *overhead* of distillation given by the number of copies of a certain primitive noisy state that are needed. As a particularly important application, we derive specific lower bounds on the overhead of magic state distillation [12], a leading proposal of fault-tolerant quantum computation [2,28,29]. The consequent limitations to the common distillation schemes based on quantum error correction are discussed in relation to key advances in the search for better codes [30–33]. Lastly, we provide a no-go theorem for the simulation of unitary resource channels, which is analogous to state purification, in accordance with the recent interest in extending conventional resource theory approaches for quantum states to quantum channels (see, e.g., [34–41] for general treatments).

We start by introducing the notations. The sets of free operations and free states are respectively denoted by $\mathcal{O}$ and $\mathcal{F}$. They obey the golden rule that $\mathcal{O} \subseteq \tilde{\mathcal{O}}$, where $\tilde{\mathcal{O}} \coloneqq \{\mathcal{E} \mid \forall\, \rho \in \mathcal{F}, \mathcal{E}(\rho) \in \mathcal{F}\}$ (commonly known as the set of resource nongenerating operations in the literature). Note that virtually no assumptions on the specific properties of the resource theory are needed in this work, that is, $\mathcal{F}$ is almost completely up to one's choice as long as there exists some resource pure state (technically, $\mathcal{F}$ is topologically closed and $\exists \psi \notin \mathcal{F}$) so that the purification task is well-defined. Even the convexity of $\mathcal{F}$, which is a common postulate for general resource theory results and frameworks (see, e.g., [22,36,42–45]), is not needed.

The general goal of purification tasks is to transform some noisy primitive state to a pure target resource state by some protocol represented by a free operation. In this work, we make a mild assumption that the density matrix representing the primitive state is full rank, which holds generically for common noise effects and settings of practical interest such as multiple noisy qubits. We would also want to consider protocols that produce desired outputs with a certain probability as long as we know when they do so (an important example being magic state distillation, as we shall discuss later). To encompass such cases, consider the generalization of $\tilde{\mathcal{O}}$ to the class $\tilde{\mathcal{O}}_{\rm sub} \coloneqq \{\mathcal{L} \mid \forall\, \rho \in \mathcal{F}, \exists\, t \geq 0, \sigma \in \mathcal{F}, \text{s.t.} \mathcal{L}(\rho) = t \cdot \sigma\}$, which consists of subnormalized quantum operations (suboperations), i.e., completely positive and trace-nonincreasing maps. A free probabilistic protocol that transforms ρ to γ with probability p and accuracy $1-\epsilon$ (or error ϵ) is modeled by a quantum operation $\mathcal{E}_{A\to XB}$ such that $\mathcal{E}_{A\to XB}(\rho_A) = |0\rangle\langle 0|_X \otimes \mathcal{L}_{A\to B}(\rho_A) + |1\rangle\langle 1|_X \otimes \mathcal{G}_{A\to B}(\rho_A)$. Here X is an external flag register that keeps track of whether the protocol succeeds (0) or not (1); $\mathcal{L} \in \mathcal{O}_{\rm sub}$ (any $\mathcal{O}_{\rm sub} \subseteq \tilde{\mathcal{O}}_{\rm sub}$) is the free suboperation representing the successful transformation such that $\mathcal{L}_{A\to B}(\rho_A) = p\tau_B$ where $p = \mathrm{Tr}\mathcal{L}(\rho)$ and τ is a density matrix satisfying $F(\tau,\gamma) \geq 1-\epsilon$ where $F(\rho,\sigma) \coloneqq \|\sqrt{\rho}\sqrt{\sigma}\|_1^2$ is the fidelity between ρ and σ. The case where $\mathcal{L}$ is a completely positive trace preserving (CPTP) map and thus $p=1$ corresponds to a deterministic protocol.

Now we are ready to introduce the explicit results. The following theorem reveals fundamental limitations on the accuracy and success probability of resource purification.

Theorem 1. Given any full-rank primitive state $\rho \notin \mathcal{F}$ and any pure target resource state $\psi \notin \mathcal{F}$, the following relation between the success probability p and transformation error ϵ must hold for any free probabilistic protocol:

$$\frac{\epsilon}{p} \geq \frac{\lambda_{\min}(\rho)(1-f_\psi)}{1+R(\rho)}, \tag{1}$$

where $\lambda_{\min}(\rho)$ is the smallest eigenvalue of ρ, $f_\psi \coloneqq \max_{\omega\in\mathcal{F}} \mathrm{Tr}(\psi\omega)$ is the maximum overlap between ψ and free states $\mathcal{F}$, and $R(\rho) \coloneqq \min\{s \mid \exists\, s \geq 0, \text{state } \sigma, \text{s.t.} (\rho + s\sigma)/(1+s) \in \mathcal{F}\}$ is the generalized robustness of state ρ. For the deterministic case ($p=1$), the bound can be improved to $\epsilon \geq \lambda_{\min}(\rho)(1-f_\psi)$.

Notice that $f_\psi < 1$ always holds by its definition, so the bound is always greater than zero, meaning that there is always a neighborhood of any ψ that cannot be reached by any free protocol. This theorem establishes an "uncertainty relation" between the accuracy and success probability of purification characterized by a regime of $\{\epsilon, p\}$ that is not achievable by any free protocol, as illustrated in Fig. 1. In particular, by letting $\epsilon = 0$, we directly rule out the possibility of perfect purification:

Corollary 1. It is impossible to exactly transform a full-rank primitive state to a pure target resource state by any free protocol, even probabilistically.

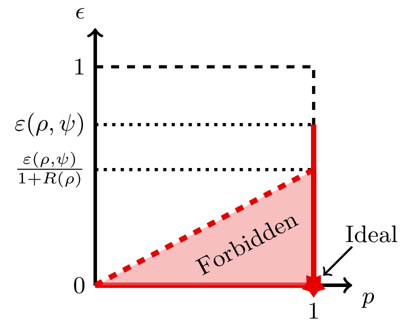

FIG. 1. Interplay between the transformation error ϵ and success probability p. The lower right corner represents the most ideal scenario where ϵ is small and p is large. The red region and solid lines represent the forbidden regime such that no purification protocol with the corresponding parameters can exist. $\varepsilon(\rho,\psi) = \lambda_{\min}(\rho)(1-f_\psi)$.

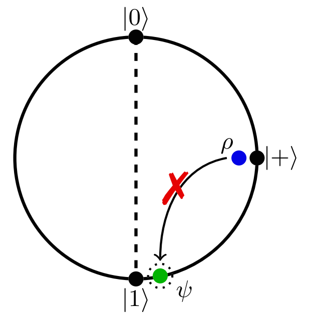

FIG. 2. A qubit coherence theory example illustrated using the Bloch sphere. Here ρ is a mixed state close to the maximally coherent state $|+\rangle$, and ψ is a pure state close to the basis state $|1\rangle$. Our no-go theorems indicate that an arbitrarily accurate probabilistic transformation from ρ to ψ is impossible.

Below we sketch our approach to proving the above results. See the Supplemental Material [46] for the detailed proof and extended discussions.

Proof. (Sketch) The cornerstone of our proof is an information-theoretic quantity called the quantum hypothesis testing relative entropy [48,52], which is defined as $D_H^\epsilon(\rho\|\sigma) := -\log\min\{\mathrm{Tr}M\sigma|\mathrm{Tr}\rho M \geq 1-\epsilon, 0 \leq M \leq \mathbb{1}\}$ for two quantum states ρ and σ. The induced resource measure given by $\mathfrak{D}_H^\epsilon(\rho) := \min_{\omega\in\mathcal{F}} D_H^\epsilon(\rho\|\omega)$, which was recently related to the rates of certain one-shot resource trading tasks [22], is shown to exhibit a peculiar property: for any full-rank ρ, it vanishes at $\epsilon = 0$ and is continuous around it. The proof then follows from suitably combining this property with the monotonicity of $\mathfrak{D}_H$ (nonincreasing under free operations). ■

Note that Ref. [53] reached a similar conclusion for time-translationally invariant operations in coherence theory. Also note that the full-rank assumption and the error bound can be improved in certain cases by different proof methods, which will be elaborated in follow-up works.

Remarkably, the noisy primitive state ρ could be much more valuable in terms of other resource measures and tasks or live in much higher dimensions than the pure target state ψ. However, the possibility of trading ρ for ψ, even probabilistically, is ruled out. This should be contrasted with the case of pure input ρ, where there are no such limitations. An illustrative toy example in terms of the theory of coherence is given in Fig. 2, where ρ is a slightly noisy version of the maximally coherent state $|+\rangle$ (which can be arbitrarily close to $|+\rangle$), while ψ is a pure target state very close to the basis (incoherent) state $|1\rangle$. It is clear from geometrical intuitions that common coherence measures (see, e.g., [11]) assign much greater value to ρ than to ψ, and it is known that $|+\rangle$ can be transformed to any other state, including ψ [18,54]. However, our results indicate that there is always a neighborhood of ψ that cannot be reached starting from ρ. This highlights the special role of $\mathfrak{D}_H$ among all resource measures and indicates sharp distinctions between pure state transformation problems and mixed state ones.

The following scheme of resource purification, usually known as "distillation" or "concentration," is of the greatest practical interest: one has access to multiple copies of some noisy primitive resource state and the goal is to "distill" certain useful pure resource states to some desired accuracy by free operations while consuming as few copies of the primitive state as possible. Most notably, the distillation of entanglement [15–17], coherence [18–20], and magic states [12] has been extensively studied as a key subroutine in quantum communication and computation. Therefore, the amount of primitive states needed to accomplish the desired distillation, namely the resource cost or overhead, is a key figure of merit for distillation protocols. To present the most general result, we consider error on the entire output state (which could be a collection of unit states) for now. As we now show, our no-go theorems indicate fundamental lower bounds on the total overhead of distillation.

Theorem 2. Consider the task of distilling some pure target resource state ψ, with error at most ϵ, from n copies of primitive state $\hat{\rho}$. For any full-rank $\hat{\rho}$, there does not exist any probabilistic protocol with success probability p that accomplishes the task if the following is not satisfied:

$$n \geq \log_{[1+R(\hat{\rho})]/\lambda_{\min}(\hat{\rho})} \frac{(1-f_\psi)p}{\epsilon}. \tag{2}$$

For deterministic case ($p = 1$), the bound can be improved to $n \geq \log_{1/\lambda_{\min}(\hat{\rho})}(1-f_\psi)/\epsilon$.

Proof. Let $\hat{\rho}^{\otimes n}$ be the primitive state in Theorem 1. Notice that $\lambda_{\min}(\hat{\rho}^{\otimes n}) = \lambda_{\min}(\hat{\rho})^n$. For the deterministic case, Theorem 1 implies that for any full-rank state $\hat{\rho}$, we have

$$\epsilon \geq \lambda_{\min}(\hat{\rho}^{\otimes n})(1-f_\psi) = \lambda_{\min}(\hat{\rho})^n(1-f_\psi). \tag{3}$$

This directly translates to $n \geq \log_{1/\lambda_{\min}(\hat{\rho})}(1-f_\psi)/\epsilon$. For the probabilistic case, note the following: by the definition of $R(\hat{\rho})$, there exists some state τ such that $\hat{\rho}+R(\hat{\rho})\tau=[1+R(\hat{\rho})]\omega$ where $\omega\in\mathcal{F}$. By expanding this equation, we obtain

$$\omega^{\otimes n}=\frac{1}{[1+R(\hat{\rho})]^n}\hat{\rho}^{\otimes n}+\frac{[1+R(\hat{\rho})]^n-1}{[1+R(\hat{\rho})]^n}\tau', \qquad (4)$$

where $\omega^{\otimes n}\in\mathcal{F}$ axiomatically [42] and τ' is a density operator. Therefore, $1+R(\hat{\rho}^{\otimes n})\leq[1+R(\hat{\rho})]^n$. Now, by Theorem 1, for any full-rank state $\hat{\rho}'$, we have

$$\epsilon/p\geq\frac{\lambda_{\min}(\hat{\rho}'^{\otimes n})(1-f_\psi)}{1+R(\hat{\rho}'^{\otimes n})}\geq\frac{\lambda_{\min}(\hat{\rho}')^n(1-f_\psi)}{[1+R(\hat{\rho}')]^n}. \qquad (5)$$

This directly translates to Eq. (2). ■

The above theorem indicates that for distillation protocols that succeed with at least a constant probability (that does not vanish when reducing the target ϵ), the total overhead must scale at least as $\Omega[\log(1/\epsilon)]$ as $\epsilon\to 0$.

As an important application, we discuss magic state distillation [12], which is a major component of the leading scheme for fault tolerance [2,28,29]. Here, the so-called Clifford operations are considered free since they admit fault-tolerant implementations thanks to stabilizer codes [2,55–57], but meanwhile their computational power is very limited—due to the celebrated Gottesman–Knill theorem, they can even be efficiently simulated by classical computers [2,58,59]. To achieve universal quantum computation, one needs non-Clifford gates such as $T=\mathrm{diag}(1,e^{i\pi/4})$. A standard approach is to distill high-quality magic state $|T\rangle=(|0\rangle+e^{i\pi/4}|1\rangle)/\sqrt{2}$ from sufficiently many noisy magic states off-line and then use an approximate $|T\rangle$ state to emulate each low-error logical T gate in the circuit via a technique called state injection or gadgetization [60]. Since the resource cost of this magic state distillation component is dominant in the entire scheme, it is crucial to understand the ultimate limitations to its efficiency.

We now address this problem by tailoring our general results to the practical magic state distillation settings, providing the first rigorous understanding of the resources required for fault-tolerance schemes. (Note that the resource theory ideas have advanced our understanding of magic states and quantum computation in various other ways [13,27,61–64].) Known protocols for magic state distillation are commonly based on concatenating error correction subroutines using stabilizer codes to probabilistically produce an output with sufficiently high quality upon passing the syndrome measurements. The output could take the form of a large state with each marginal sufficiently close to a unit target state, in which case we are also interested in the average overhead, i.e., the total overhead divided by the number of marginals. Here we only showcase the T-state result, but the bounds for other useful magic states (see, e.g., [33]) can be similarly obtained by plugging in corresponding parameters.

Theorem 3. Consider the following general form of magic state distillation task: given n copies of full-rank primitive magic states $\hat{\rho}$, output an m-qubit state τ such that $\mathrm{Tr}\tau_i T=\langle T|\tau_i|T\rangle\geq 1-\epsilon,\ \forall\ i=1,\ldots,m$ where $\tau_i=\mathrm{Tr}_{\bar{i}}\tau$ is the ith qubit. Then the average overhead of any free probabilistic protocol that succeeds with probability p must obey

$$C\coloneqq n/m\geq\frac{1}{m}\log_{[1+R(\hat{\rho})]/[\lambda_{\min}(\hat{\rho})]}\frac{[(4-2\sqrt{2})^m-1]p}{(4-2\sqrt{2})^m m\epsilon}. \qquad (6)$$

Proof. By applying the union bound, we have $\langle T^{\otimes m}|\tau|T^{\otimes m}\rangle\geq 1-m\epsilon$. Also notice that $f_{T^{\otimes m}}=(4-2\sqrt{2})^{-m}$ [22,65–67]. By plugging everything into Eq. (2), we obtain the claimed bound. ■

In the analyses of magic state distillation protocols, one is particularly interested in the exponent γ in the asymptotic average overhead $O[\log^\gamma(1/\epsilon)]$ as $\epsilon\to 0$. A subtlety of our lower bound is that the output size m could depend on the target ϵ for specific protocols. Thus, to understand the scaling, one needs to take into account the behavior of m as well. There are two key implications of our bound to code-based distillation protocols. Assuming nonvanishing success probability (the passing probability of deeper rounds of concatenation converges sufficiently fast to one), we conclude the following: (i) It is impossible to construct a protocol with sublogarithmic average overhead ($\gamma<1$) with any $[n,k,d]$ code such that $k\leq d$. This can be seen by plugging $m=k^\nu$ and $\log(1/\epsilon)\sim d^\nu$ into Eq. (6). This in particular implies a $\gamma\geq 1$ bound for $k=1$ codes in response to open questions raised in, e.g., [30,31]. Note that the best known such codes allow $\gamma\to 2$ [32,33], so there is still potential room for improvement. (ii) Any $\gamma<1$ protocol must have a scale (size of the output) that diverges under concatenation. It was actually believed that no codes allowing $\gamma<1$ exist [30], but the recent breakthrough work by Hastings and Haah [31] gives a peculiar example of such a code (see also [68]), prompting the question of whether there is any fundamental limit. (There, indeed, the codes employed have $k>d$.) Our results indicate that, although the average overhead of such a protocol is considered low, its output size must grow rapidly as we reduce ϵ, which inevitably blows up the overall cost.

Finally, we make a basic extension to the channel resource theory setting (see, e.g., [34,39,40]), a more general setting of surging interest recently, which directly applies to quantum channels, gates, and dynamical processes, etc. We show that, under the analogous golden rule, it is generally impossible to perfectly transform a noisy quantum channel into a unitary resource channel. A straightforward implication of this result is that the zero-error quantum capacity of generic noisy channels, e.g., the depolarizing channel, is zero. See the Supplemental Material [46] for detailed statements and proofs. More comprehensive studies of the channel setting will be left for follow-up.

To conclude, this work establishes quantitative bounds on the accuracy and efficiency of purifying noisy quantum resources and thus draws practical boundaries for quantum error correction and mitigation, by employing one-shot quantum resource theory techniques. Our results universally apply to quantum resources of any reasonable kind. The bounds depend only on very few parameters that concisely encode relevant properties of the noise, the target state, and the resource theory and are thus easy to analyze. Like the no-cloning theorem, our "no-purification" theorems stem from fundamental laws of quantum mechanics at bottom. We demonstrate the power and practical relevance of our general methods by establishing strong lower bounds on the overhead of distillation tasks (e.g., magic state distillation), which provide rigorous understandings of and useful benchmarks for the resource requirements of practical quantum technologies, in particular fault-tolerant quantum computation, as the Heisenberg limit did for quantum metrology.

An important future work is to investigate to what extent our various bounds can be approached, both by general means and in specific theories. For instance, it remains to be checked how close the state-of-the-art entanglement purification protocols (see, e.g., [69]) are to the fundamental limits set here. We also expect our general, primary results to see improvements in various cases and, more generally, stimulate further studies on optimal quantum resource purification. It would also be interesting to further understand the approximate and probabilistic regimes of unitary channel simulation due to its connections to the fields of quantum Shannon theory, gate and circuit synthesis, etc. In sum, a key message of this work is that the cost of practically implementing quantum technologies or experiments could not be indefinitely improved in general due to noise effects. As we are now witnessing an exciting paradigm shift from blueprinting quantum advantages in theory to actually putting them into practice [14,70], we anticipate that such a rigorous understanding of the fundamental obstacles will serve as an important guideline and have far-reaching implications for quantum science and technology.

We thank Earl Campbell, Daniel Gottesman, Gilad Gour, Jeongwan Haah, Sirui Lu, Bartosz Regula, Ryuji Takagi, Julio I. de Vicente, and Andreas Winter for helpful discussions and feedback and the anonymous referees for several valuable comments. K. F. is supported by the University of Cambridge Isaac Newton Trust Early Career Grant No. RG74916. Z.-W. L. is supported by the Perimeter Institute for Theoretical Physics. Research at the Perimeter Institute is supported by the government of Canada through Industry Canada and by the province of Ontario through the Ministry of Research and Innovation.

*Corresponding author.
kf383@cam.ac.uk
†Corresponding author.
zliu1@perimeterinstitute.ca

[1] J. P. Dowling and G. J. Milburn, Quantum technology: The second quantum revolution, Phil. Trans. R. Soc. A **361**, 1655 (2003).
[2] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information: 10th Anniversary Edition*, 10th ed. (Cambridge University Press, New York, 2011).
[3] W. K. Wootters and W. H. Zurek, A single quantum cannot be cloned, Nature (London) **299**, 802 (1982).
[4] D. Dieks, Communication by EPR devices, Phys. Lett. **92A**, 271 (1982).
[5] P. W. Shor, Scheme for reducing decoherence in quantum computer memory, Phys. Rev. A **52**, R2493 (1995).
[6] A. M. Steane, Error Correcting Codes in Quantum Theory, Phys. Rev. Lett. **77**, 793 (1996).
[7] C. H. Bennett and G. Brassard, Quantum cryptography: Public key distribution and coin tossing, in *Proceedings of IEEE International Conference on Computers, Systems & Signal Processing, Bangalore, India* (1984), pp. 175–179.
[8] V. Scarani, S. Iblisdir, N. Gisin, and A. Acín, Quantum cloning, Rev. Mod. Phys. **77**, 1225 (2005).
[9] B. Wootters and W. Zurek, The no-cloning theorem, Phys. Today **62**, 2, 76 (2009).
[10] R. Horodecki, P. Horodecki, M. Horodecki, and K. Horodecki, Quantum entanglement, Rev. Mod. Phys. **81**, 865 (2009).
[11] A. Streltsov, G. Adesso, and M. B. Plenio, Colloquium: Quantum coherence as a resource, Rev. Mod. Phys. **89**, 041003 (2017).
[12] S. Bravyi and A. Kitaev, Universal quantum computation with ideal Clifford gates and noisy ancillas, Phys. Rev. A **71**, 022316 (2005).
[13] V. Veitch, S. A. H. Mousavian, D. Gottesman, and J. Emerson, The resource theory of stabilizer quantum computation, New J. Phys. **16**, 013009 (2014).
[14] J. Preskill, Quantum Computing in the NISQ era and beyond, Quantum **2**, 79 (2018).
[15] C. H. Bennett, H. J. Bernstein, S. Popescu, and B. Schumacher, Concentrating partial entanglement by local operations, Phys. Rev. A **53**, 2046 (1996).
[16] C. H. Bennett, G. Brassard, S. Popescu, B. Schumacher, J. A. Smolin, and W. K. Wootters, Purification of Noisy Entanglement and Faithful Teleportation via Noisy Channels, Phys. Rev. Lett. **76**, 722 (1996).
[17] C. H. Bennett, D. P. DiVincenzo, J. A. Smolin, and W. K. Wootters, Mixed-state entanglement and quantum error correction, Phys. Rev. A **54**, 3824 (1996).
[18] A. Winter and D. Yang, Operational Resource Theory of Coherence, Phys. Rev. Lett. **116**, 120404 (2016).
[19] K. Fang, X. Wang, L. Lami, B. Regula, and G. Adesso, Probabilistic Distillation of Quantum Coherence, Phys. Rev. Lett. **121**, 070404 (2018).
[20] B. Regula, K. Fang, X. Wang, and G. Adesso, One-Shot Coherence Distillation, Phys. Rev. Lett. **121**, 010401 (2018).
[21] E. Chitambar and G. Gour, Quantum resource theories, Rev. Mod. Phys. **91**, 025001 (2019).
[22] Z.-W. Liu, K. Bu, and R. Takagi, One-Shot Operational Quantum Resource Theory, Phys. Rev. Lett. **123**, 020401 (2019).

[23] F. G. S. L. Brandao and N. Datta, One-shot rates for entanglement manipulation under non-entangling maps, IEEE Trans. Inf. Theory **57**, 1754 (2011).
[24] K. Fang, X. Wang, M. Tomamichel, and R. Duan, Non-asymptotic entanglement distillation, IEEE Trans. Inf. Theory **65**, 6454 (2019).
[25] Q. Zhao, Y. Liu, X. Yuan, E. Chitambar, and A. Winter, One-shot coherence distillation: Towards completing the picture, IEEE Trans. Inf. Theory **65**, 6441 (2019).
[26] N. Yunger Halpern and J. M. Renes, Beyond heat baths: Generalized resource theories for small-scale thermodynamics, Phys. Rev. E **93**, 022126 (2016).
[27] X. Wang, M. M. Wilde, and Y. Su, Efficiently Computable Bounds for Magic State Distillation, Phys. Rev. Lett. **124**, 090505 (2020).
[28] P. W. Shor, Fault-tolerant quantum computation, in *Proceedings of 37th Conference on Foundations of Computer Science* (1996), pp. 56–65.
[29] E. T. Campbell, B. M. Terhal, and C. Vuillot, Roads towards fault-tolerant universal quantum computation, Nature (London) **549**, 172 (2017).
[30] S. Bravyi and J. Haah, Magic-state distillation with low overhead, Phys. Rev. A **86**, 052329 (2012).
[31] M. B. Hastings and J. Haah, Distillation with Sublogarithmic Overhead, Phys. Rev. Lett. **120**, 050504 (2018).
[32] J. Haah, M. B. Hastings, D. Poulin, and D. Wecker, Magic state distillation with low space overhead and optimal asymptotic input count, Quantum **1**, 31 (2017).
[33] J. Haah and M. B. Hastings, Codes and Protocols for Distilling T, controlled-S, and Toffoli Gates, Quantum **2**, 71 (2018).
[34] G. Gour, Comparison of quantum channels by superchannels, IEEE Trans. Inf. Theory **65**, 5880 (2019).
[35] M. M. Wilde, Entanglement cost and quantum channel simulation, Phys. Rev. A **98**, 042338 (2018).
[36] L. Li, K. Bu, and Z.-W. Liu, Quantifying the resource content of quantum channels: An operational approach, Phys. Rev. A **101**, 022335 (2020).
[37] T. Theurer, D. Egloff, L. Zhang, and M. B. Plenio, Quantifying Operations with an Application to Coherence, Phys. Rev. Lett. **122**, 190405 (2019).
[38] R. Takagi and B. Regula, General Resource Theories in Quantum Mechanics and Beyond: Operational Characterization via Discrimination Tasks, Phys. Rev. X **9**, 031053 (2019).
[39] Y. Liu and X. Yuan, Operational resource theory of quantum channels, Phys. Rev. Research **2**, 012035 (2020).
[40] Z.-W. Liu and A. Winter, Resource theories of quantum channels and the universal role of resource erasure, arXiv:1904.04201.
[41] G. Gour and A. Winter, How to Quantify a Dynamical Quantum Resource, Phys. Rev. Lett. **123**, 150401 (2019).
[42] F. G. S. L. Brandão and G. Gour, Reversible Framework for Quantum Resource Theories, Phys. Rev. Lett. **115**, 070503 (2015).
[43] Z.-W. Liu, X. Hu, and S. Lloyd, Resource Destroying Maps, Phys. Rev. Lett. **118**, 060502 (2017).
[44] R. Takagi, B. Regula, K. Bu, Z.-W. Liu, and G. Adesso, Operational Advantage of Quantum Resources in Subchannel Discrimination, Phys. Rev. Lett. **122**, 140402 (2019).
[45] B. Regula, K. Bu, R. Takagi, and Z.-W. Liu, Benchmarking one-shot distillation in general quantum resource theories, Phys. Rev. A **101**, 062315 (2020).
[46] See Supplemental Material, which includes Refs. [19,40,43,47–51], at http://link.aps.org/supplemental/10.1103/PhysRevLett.125.060405 for technical derivations and proofs of the results in this Letter.
[47] R. A. Horn and C. R. Johnson, *Matrix Analysis* (Cambridge University Press, England, 2012).
[48] L. Wang and R. Renner, One-Shot Classical-Quantum Capacity and Hypothesis Testing, Phys. Rev. Lett. **108**, 200501 (2012).
[49] G. Chiribella, G. M. D'Ariano, and P. Perinotti, Transforming quantum operations: Quantum supermaps, Europhys. Lett. **83**, 30004 (2008).
[50] K. Fang, X. Wang, M. Tomamichel, and M. Berta, Quantum Channel Simulation and the Channel's Smooth Max-Information, in *2018 IEEE International Symposium on Information Theory (ISIT)* (IEEE, New York, 2018), pp. 2326–2330.
[51] K. Fang, X. Wang, M. Tomamichel, and M. Berta, Quantum Channel Simulation and the Channel's Smooth Max-Information, in IEEE Trans. Inf. Theory **66**, 2129 (2020).
[52] F. Buscemi and N. Datta, The quantum capacity of channels with arbitrarily correlated noise, IEEE Trans. Inf. Theory **56**, 1447 (2010).
[53] I. Marvian, Coherence distillation machines are impossible in quantum thermodynamics, Nat. Commun. **11**, 25 (2020).
[54] T. Baumgratz, M. Cramer, and M. B. Plenio, Quantifying Coherence, Phys. Rev. Lett. **113**, 140401 (2014).
[55] D. Gottesman, Class of quantum error-correcting codes saturating the quantum Hamming bound, Phys. Rev. A **54**, 1862 (1996).
[56] D. Gottesman, Stabilizer codes and quantum error correction, Ph.D. thesis, California Institute of Technology (1997).
[57] A. R. Calderbank, E. M. Rains, P. W. Shor, and N. J. A. Sloane, Quantum Error Correction and Orthogonal Geometry, Phys. Rev. Lett. **78**, 405 (1997).
[58] D. Gottesman, The Heisenberg Representation of Quantum Computers, arXiv:quant-ph/9807006.
[59] S. Aaronson and D. Gottesman, Improved simulation of stabilizer circuits, Phys. Rev. A **70**, 052328 (2004).
[60] D. Gottesman and I. L. Chuang, Demonstrating the viability of universal quantum computation using teleportation and single-qubit operations, Nature (London) **402**, 390 (1999).
[61] M. Howard and E. Campbell, Application of a Resource Theory for Magic States to Fault-Tolerant Quantum Computing, Phys. Rev. Lett. **118**, 090501 (2017).
[62] J. R. Seddon and E. T. Campbell, Quantifying magic for multi-qubit operations, Proc. R. Soc. A **475**, 20190251 (2019).
[63] X. Wang, M. M. Wilde, and Y. Su, Quantifying the magic of quantum channels, New J. Phys. **21**, 103002 (2019).
[64] M. Beverland, E. Campbell, M. Howard, and V. Kliuchnikov, Lower bounds on the non-Clifford resources for quantum computations, arXiv:1904.01124.
[65] E. T. Campbell, Catalysis and activation of magic states in fault-tolerant architectures, Phys. Rev. A **83**, 032317 (2011).

[66] S. Bravyi and D. Gosset, Improved Classical Simulation of Quantum Circuits Dominated by Clifford Gates, Phys. Rev. Lett. **116**, 250501 (2016).
[67] S. Bravyi, D. Browne, P. Calpin, E. Campbell, D. Gosset, and M. Howard, Simulation of quantum circuits by low-rank stabilizer decompositions, Quantum **3**, 181 (2019).
[68] A. Krishna and J.-P. Tillich, Towards Low Overhead Magic State Distillation, Phys. Rev. Lett. **123**, 070507 (2019).
[69] S. Krastanov, V. V. Albert, and L. Jiang, Optimized Entanglement Purification, Quantum **3**, 123 (2019).
[70] F. Arute *et al.*, Quantum supremacy using a programmable superconducting processor, Nature (London) **574**, 505 (2019).