

Quantifying Quantum Speedups: Improved Classical Simulation From Tighter Magic Monotones

James R. Seddon^{1,*}, Bartosz Regula^{2,†}, Hakop Pashayan^{3,4,5}, Yingkai Ouyang⁶, and Earl T. Campbell^{6,7,‡}

¹*Department of Physics and Astronomy, University College London, London, United Kingdom*

²*School of Physical and Mathematical Sciences, Nanyang Technological University, 637371, Singapore*


³*Institute for Quantum Computing and Department of Combinatorics and Optimization, University of Waterloo, Ontario N2L 3G1, Canada*

⁴*Perimeter Institute for Theoretical Physics, Waterloo, Ontario N2L 2Y5 Canada*

⁵*Centre for Engineered Quantum Systems, School of Physics, The University of Sydney, Sydney, New South Wales 2006, Australia*

⁶*Department of Physics & Astronomy, University of Sheffield, Sheffield S3 7RH, United Kingdom*

⁷*AWS Center for Quantum Computing, Pasadena, California 91125, USA*

 (Received 13 March 2020; revised 11 September 2020; accepted 22 February 2021; published 18 March 2021)

Consumption of magic states promotes the stabilizer model of computation to universal quantum computation. Here, we propose three different classical algorithms for simulating such universal quantum circuits, and characterize them by establishing precise connections with a family of magic monotones. Our first simulator introduces a new class of quasiprobability distributions and connects its runtime to a generalized notion of negativity. We prove that this algorithm has significantly improved exponential scaling compared to all prior quasiprobability simulators for qubits. Our second simulator is a new variant of the stabilizer-rank simulation algorithm, extended to work with mixed states and with significantly improved runtime bounds. Our third simulator trades precision for speed by discarding negative quasiprobabilities. We connect each algorithm's performance to a corresponding magic monotone and, by comprehensively characterizing the monotones, we obtain a precise understanding of the simulation runtime and error bounds. Our analysis reveals a deep connection between all three seemingly unrelated simulation techniques and their associated monotones. For tensor products of single-qubit states, we prove that our monotones are all equal to each other, multiplicative and efficiently computable, allowing us to make clear-cut comparisons of the simulators' performance scaling. Furthermore, our monotones establish several asymptotic and nonasymptotic bounds on state interconversion and distillation rates. Beyond the theory of magic states, our classical simulators can be adapted to other resource theories under certain axioms, which we demonstrate through an explicit application to the theory of quantum coherence.

DOI: [10.1103/PRXQuantum.2.010345](https://doi.org/10.1103/PRXQuantum.2.010345)

I. INTRODUCTION

Classical simulation of quantum systems has a long and fruitful history. Insurmountable obstructions to the classical simulation of quantum systems gave birth to the field of quantum computation [1] and the search for quantum computational resources. Despite the computational limitations

of classical simulation, surprisingly powerful classical simulators have since been discovered including simulators of stabilizer circuits [2,3], fermionic linear optics and matchgates [4–7], and others [8–13]. Improvement and characterization of classical simulation algorithms helps benchmark the computational speedups that quantum computers can provide and also provides tools useful in their own right [14–16].

Stabilizer circuits are initialized in so-called stabilizer states and evolved by stabilizer operations, such that the system stays in a stabilizer state throughout the whole computation. These circuits are important in fault-tolerant quantum computation and can be efficiently classically simulated by virtue of the Gottesman-Knill theorem [2]. An elegant extension of stabilizer circuits enables them to perform universal computation by allowing the input states

*jseddonquantum@gmail.com

†bartosz.regula@gmail.com

‡earlrcampbell@gmail.com

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

to include so-called *magic states* [17,18]. Aaronson and Gottesman [3] showed how to classically simulate such circuits with a runtime that scales exponentially with the number of input magic state qubits, yet still scales efficiently with respect to the number of stabilizer-state qubits. Consequently, we can perform an efficient classical simulation for any class of circuits that is *nearly stabilizer* in the sense that they use only logarithmically many input magic state qubits. Subsequent developments showed that the difficulty of simulating a quantum circuit depends not only on the number of magic state inputs, but also on the type of magic that these states possess.

In the pursuit of faster classical simulation of nearly stabilizer circuits, two leading approaches have emerged: quasiprobability [15,19–22] and stabilizer rank–based [23–26] simulators. These simulators all have their runtime determined by a function called a *magic monotone* that quantifies how far the magic states deviate from the set of stabilizer states. With these modern simulators, even a very large number of magic state inputs is classically tractable, provided the magic states are close enough to stabilizer states, as quantified by the relevant magic monotone. However, different simulators come with their own magic monotone and therefore different runtime scalings. So far, no overarching study has precisely compared the runtimes and monotones for different stabilizer simulators. The difficulty of comparison is exacerbated since some monotones are not easily calculated. We next review these simulation methods, before stating our main results that further sharpen the performance of modern simulators and reveal a cohesive picture of a previously fragmented landscape of simulators.

A. Review of prior art

Quasiprobability simulators work by representing the target quantum state by an operator probabilistically chosen from a discrete set known as a *frame* [20,27]. Examples of relevant frames include the set of density operators corresponding to pure stabilizer states [21], the set of Pauli operators [28], and the set of phase-point operators [20] used in the construction of the discrete Wigner function [29,30]. Importantly, given a choice of a classically simulable frame, any input state, which is a convex combination of frame elements admits an efficient classical simulation algorithm [31]. In Ref. [20], Pashayan *et al.* showed that when the input state is a nonconvex linear combination of frame elements, the only source of inefficiency in the runtime of quasiprobabilistic algorithms is given by the *negativity* of the state—a frame-dependent quantity, which measures the degree of departure from convex mixtures of frame elements.

Quantum systems consisting of odd-dimensional subsystems (qudits) [32–34] admit an especially natural choice of frame. Here, the frame can be fixed to a set of

phase-point operators for which the convex combinations of frame elements are the states with a positive discrete Wigner function [29,30]. All qudit stabilizer states have a positive Wigner function, which leads to efficient, classical simulation of qudit stabilizer circuits [31]. The negativity under this choice of frame was shown in Ref. [20] to correspond to the *mana* $\mathcal{M}(\cdot)$ —a magic monotone introduced in Ref. [35]. Notably, the mana has the convenient property that it is multiplicative [36], i.e., $\mathcal{M}(\rho \otimes \sigma) = \mathcal{M}(\rho)\mathcal{M}(\sigma)$. Computations of the mana in large dimensions are generally extremely difficult, but, due to multiplicativity, they are significantly simplified for products of states on smaller systems. Multiplicativity of operationally meaningful monotones allows for an easy evaluation of related quantities, such as a simulator’s runtime or bounds on asymptotic rates of state conversion.

Curiously, for the fundamentally important case of qubits, phase-point operator frames do not possess many of the aforementioned desirable properties. A straightforward application of techniques that work for qudits yields a Wigner function that can be negative for some pure stabilizer states. Although alternative ways of defining a well-behaved Wigner function for qubits are possible, they always [37] suffer from drawbacks such as the free operations and states being restricted to a subclass of the usual free operations [38,39]; or the monotones being supermultiplicative and the set of positively represented states not being closed under tensor product [40]. Quasiprobability simulators based on qubit phase-point operator frames inherit these limitations, prompting alternative approaches.

In Ref. [21] Howard and Campbell presented a quasiprobability simulator for qubits using a frame composed of projectors onto pure stabilizer states. They showed that this gives rise to a classical simulation algorithm with a runtime linked to a magic monotone called the *robustness of magic*. It is a qubit-based simulator that permits and utilizes the simulation of noisy inputs and operations, and possesses many desirable traits. However, presently, quasiprobability simulators are slower than stabilizer-rank simulators; additionally, the robustness of magic is nonmultiplicative and extremely difficult to compute, even in the asymptotic regime for products of relevant single-qubit states [17,41].

A seemingly independent line of work on classical simulation was introduced in Ref. [23] with the stabilizer rank–based simulators [23,25,42–44]. These simulators achieve a stronger notion of simulation [45] by approximately sampling from the output distribution of the quantum circuit. However, they can only simulate pure states and operations and have not previously been generalized to noisy quantum circuits. Stabilizer-rank simulators represent the initial quantum-state vector as a superposition of stabilizer states, and the only source of inefficiency in runtime is introduced by the exponential number of terms required to represent states in this way—the minimal

number of such terms being precisely the stabilizer rank. The original algorithm had a runtime quadratic in the stabilizer rank, but this was later improved by the development of fast norm estimation [24] that provides a runtime linear in stabilizer rank. This has resulted in a sizable runtime advantage for stabilizer-rank simulators, and it is currently unclear if a similar improvement in quasiprobabilistic methods is possible. To circumvent the difficulty in computing the stabilizer rank as well as its nonmultiplicative behavior, Ref. [24] also introduced the notion of approximate stabilizer rank, which was later related to a monotone called the *stabilizer extent* [25]. While the extent is in general not multiplicative [46], it is multiplicative on any tensor product of one-, two-, and three-qubit states [25]. Unfortunately, these concepts only apply to pure states, and no mixed-state simulation method based on the stabilizer rank has been devised thus far.

B. Summary of results


In this paper, we present three new classical simulation algorithms (overviewed in Table I), which we call the *dyadic frame*, the *density-operator stabilizer rank* and the *constrained path* simulators. The algorithms allow classical simulation of general noisy stabilizer circuits with mixed magic state inputs, providing a significant extension of the capabilities of previous approaches, and

revealing connections between stabilizer rank-based and quasiprobability-based simulation. The dyadic frame and constrained path simulators produce additive precision estimates of Born-rule probabilities and Pauli observables, while the density-operator stabilizer-rank simulator approximately samples from the quantum circuit’s measurement outcome distribution. Our first two simulators trade-off quantum computational resources for additional runtime of classical simulation. The constrained path simulator, on the other hand, is always efficient in runtime, instead reducing in accuracy as the simulated quantum circuits increase in magic.


Our dyadic frame simulator is a new state-of-the-art quasiprobability simulator for qubits. Instead of sampling from stabilizer states or phase-point operators, we sample from objects we call stabilizer dyads. We show the corresponding resource monotone is smaller than the robustness of magic, leading to faster simulation runtimes. This can lead to a significantly improved exponent in the exponential scaling of the simulator’s performance: for instance, for n copies of a T state the dyadic simulator has a runtime $\mathcal{O}(4^{0.228443n})$, whereas for simulators based on the robustness of magic [21,22] the runtime is lower bounded by $\Omega(4^{0.271553n})$.

Our stabilizer-rank simulator is a new state-of-the-art simulator for sampling from qubit-based quantum circuits


TABLE I. A summary of the properties of our three classical simulation algorithms and their connections with magic monotones. Here, p_{fail} denotes the failure probability of the associated algorithm.



Dyadic frame simulator



Density-operator stabilizer-rank simulator



Constrained path simulator

Classical Simulator	Output	Associated Monotone	Method	Runtime/(poly(n) log(p_{fail}^{-1}))
Dyadic frame simulator (Sec. VII A)	Estimate of a single Born rule probability (or Pauli observable) within additive error ϵ , for arbitrary $\epsilon > 0$.	Λ Dyadic negativity	Sample dyads from quasi-probability distribution.	$\mathcal{O}(\Lambda(\rho)^2 \epsilon^{-2})$
Density-operator stabilizer rank simulator (Sec. VII B)	A bit string sampled from a distribution δ -close in the 1-norm to the quantum distribution, for arbitrary $\delta > 0$.	Ξ Mixed-state extent	Samples pure states from an ensemble then sparsifies their stabilizer rank representation.	$\mathcal{O}(\Xi(\rho)\delta^{-3})$ when $\delta \geq \delta_c$, $\mathcal{O}(\Xi(\rho)\delta^{-4})$ when $\delta < \delta_c$
Constrained path simulator (Sec. VII C)	Estimate of a single Born rule probability (or Pauli observable), with estimation error $\Delta = \mathcal{O}((\Lambda^+(\rho) - 1))$.	Λ^+ Generalized robustness	Approximates a mixed magic state ρ with a single mixed stabilizer state σ .	$\mathcal{O}(1)$

with three key technical contributions. First, our work generalizes the stabilizer-rank simulator of Refs. [24,25] from pure states to general mixed states. This allows our classical simulator to operate in and be directly comparable to more experimentally relevant regimes, where the input magic states are noisy. The natural generalization to mixed states produces a simulator with a probabilistic runtime. Second, we show this runtime can be made deterministic for an important subset of magic states. Third, we substantially improve runtimes by exploiting tighter proof techniques available in the density operator picture. Remarkably, this density operator technique is applicable and advantageous even when simulating pure states.

We show that each of our simulators—the dyadic frame, the density-operator stabilizer rank, and the constrained path simulator—is associated with a particular magic monotone, which we call the *dyadic negativity*, the *mixed-state extent*, and the *generalized robustness*, respectively. Specifically, we show that the runtime (in the case of the dyadic frame and density-operator stabilizer-rank simulators) or the precision (in the case of the constrained path simulator) of the algorithms directly relates to the corresponding magic monotone. This identifies the exponential growth of magic as the only source of inefficiency in these simulators. Crucially, we completely characterize these monotones for single-qubit states and tensor products thereof, where we prove the unexpectedly strong result that these monotones are all equal and act multiplicatively. The multiplicativity of the monotones is the first result of this type for general qubit magic monotones, and the equality between all three monotones reveals a deep and precisely quantified connection between the runtimes of stabilizer-rank and quasiprobability simulators. To the best of our knowledge, no previous work has established a quantitative connection between these, *a priori* very different, classes of simulators. All of the monotones reduce to the stabilizer extent for pure states, and so they can all be considered as generalizations of the extent to mixed states. In addition to serving as an important contribution to magic theory and tightly characterizing the resource consumption of our simulators, we use the monotones to introduce computable bounds on the asymptotic and nonasymptotic rates for magic state distillation. For some example distillation tasks, we compare our bounds to other recent results [47] and find they are much tighter across a wide parameter regime.

Classical simulation of quantum systems has been studied within various contexts other than magic theory [6,20,31,48,49], but to our knowledge none of these approaches have been adapted to the umbrella of quantum resource theories [50]. We provide a comprehensive recipe to apply our methodology to general quantum resources. We thus establish connections between a family of resource monotones and simulation tasks, shedding light on classical simulation algorithms in broader settings. For instance,

in the resource theory of quantum coherence [51,52], the ℓ_1 norm of coherence is a fundamental quantifier of this resource but lacks an operational meaning. Our results fill this gap by showing that the ℓ_1 norm of coherence quantifies the runtime of classical simulation within this theory.

This paper is structured as follows. In Sec. II we introduce the setting of magic theory and our family of monotones. Section II also discusses how the monotones connect with our simulation algorithms, providing a statement of our main theorems. In Sec. III we present a complete picture of how our monotones compare for single-qubit states by showing that they are all equal. The equality between monotones is then extended in Sec. IV to tensor-product states, where we show that the monotones are strongly multiplicative. In Sec. V we compare our new monotones with the robustness of magic and show that they can be exponentially smaller in magnitude. In Sec. VI, we discuss how the monotones can be used to bound the performance of magic state distillation protocols. Section VII contains a complete discussion of our simulation algorithms—we focus on providing an intuitive picture through illustrative examples and sketches of the main proof ideas, with the full technical details deferred to the appendix. We conclude in Sec. VIII with a discussion of our underlying assumptions and the extension of our results to resource theories beyond magic.

II. PRELIMINARIES

A. The stabilizer formalism

Here, we briefly review the stabilizer formalism. The single-qubit Pauli group \mathcal{P}_1 contains the identity matrix $\mathbb{1}$, the Pauli spin matrices X , Y , and Z , as well as their products with $\pm i\mathbb{1}$. We say a pure, single-qubit state is a stabilizer state if there exists a Pauli operator $P \in \mathcal{P}_1$, such that $P|\psi\rangle = |\psi\rangle$ and $P \neq \mathbb{1}$. There are six such states:

$$\begin{aligned} Z|0\rangle &= |0\rangle, \\ (-Z)|1\rangle &= |1\rangle, \\ \pm X|\pm\rangle &= |\pm\rangle := (|0\rangle \pm |1\rangle)/\sqrt{2}, \\ \pm Y|\pm_i\rangle &= |\pm_i\rangle := (|0\rangle \pm i|1\rangle)/\sqrt{2}. \end{aligned} \tag{1}$$

The n -qubit Pauli group \mathcal{P}_n is the group generated by tensor products of n Pauli operators. We say a pure, n -qubit state is a stabilizer state if there exists an Abelian subgroup of the Pauli group $\mathcal{S} \subset \mathcal{P}_n$ containing 2^n elements such that $S|\psi\rangle = |\psi\rangle$ for all $S \in \mathcal{S}$. The group \mathcal{S} is called the stabilizer group of the state $|\psi\rangle$, and that \mathcal{S} can be described using $\mathcal{O}(n^2)$ bits, underpinning the efficient classical simulation results of the Gottesman-Knill theorem. We use S_n to denote the set of pure n -qubit stabilizer states. The set of mixed stabilizer states \bar{S}_n is then formed by all states,

which can be decomposed as a mixture of pure stabilizers, that is, $\bar{S}_n = \text{conv}\{|\phi\rangle\langle\phi| : |\phi\rangle \in S_n\}$.

An n -qubit unitary C is Clifford if for every Pauli $P \in \mathcal{P}_n$, it follows that $CPC^\dagger \in \mathcal{P}_n$. We see that stabilizer states are mapped to stabilizer states under Clifford unitaries, and furthermore this update can be tracked efficiently. In addition, measurements of Pauli operators on stabilizer states can also be efficiently simulated by appropriately updating the stabilizer group.

We refer to the stabilizer operations as any sequence of the following: preparation of stabilizer states, Clifford unitaries, Pauli measurements, and adaptive feedforward depending on previous measurement outcomes or random coin tosses. From the perspective of complexity theory, a small caveat is required that adaptive feedforward decisions are computed using only a small (constant size) classical computer.

A quantum channel \mathcal{E} is said to be stabilizer preserving if it maps every mixed stabilizer state $\rho \in \bar{S}_n$ to another mixed stabilizer state, so $\mathcal{E}(\rho) \in \bar{S}_n$. Although meaningful when acting on the whole system in consideration, such maps can exhibit undesirable properties when acting on a part of a larger system [22]. We thus consider a relevant class of free operations defined as follows.

Definition 1. We define the set of free operations \mathcal{O}_n as the set of channels \mathcal{E} that are (i) completely positive; (ii) trace preserving, so that $\text{Tr}[\mathcal{E}(\rho)] = \text{Tr}[\rho]$; (iii) completely stabilizer preserving, in the sense that

$$[\mathcal{E} \otimes \mathbb{1}](\rho) \in \bar{S}_{2n} \quad \forall \rho \in \bar{S}_{2n}. \quad (2)$$

This set can equivalently be defined via the Choi-Jamiołkowski isomorphism as was shown in Ref. [22, Theorem 3.1].

While it is clear that the stabilizer operations are contained in \mathcal{O}_n , it is not known whether all elements of \mathcal{O}_n can be realized by the standard stabilizer operations without postselection. The Gottesman-Knill theorem has long been known to show that stabilizer operations can be efficiently classically simulated, but only recently was it shown that the more general class \mathcal{O}_n also admits efficient simulation algorithms [22]. Furthermore, it is known that certain stabilizer-preserving but non-trace-preserving maps, such as postselection on the outcome of a Pauli measurement, can also be efficiently simulated. For technical reasons we do not consider these as elements of the convex set of free operations \mathcal{O}_n in our resource theory, but we exploit their simulability in Sec. VII.

While the stabilizer operations (or the free operations \mathcal{O}_n) are not universal for quantum computation, they can be promoted to universality given an unlimited supply of a suitable nonstabilizer operation. For instance, adding the T

gate (also called the $\pi/8$ phase gate)

$$T = \begin{pmatrix} e^{i\pi/8} & 0 \\ 0 & e^{-i\pi/8} \end{pmatrix}, \quad (3)$$

promotes the stabilizer operations to full quantum universality [53]. Alternatively, one can add a supply of nonstabilizer states such as the so-called magic states:

$$|H\rangle\langle H| = (1/2) \left[\mathbb{1} + (X + Z)/\sqrt{2} \right], \quad (4)$$

$$|T\rangle\langle T| = (1/2) \left[\mathbb{1} + (X + Y)/\sqrt{2} \right], \quad (5)$$

$$|F\rangle\langle F| = (1/2) \left[\mathbb{1} + (X + Y + Z)/\sqrt{3} \right], \quad (6)$$

which we use throughout. Given a single copy of the Hadamard eigenstate $|H\rangle$ or the Clifford equivalent T state $|T\rangle$, we can perform a deterministic T gate using state injection [17]. Therefore, full universality can be achieved given stabilizer operations and a supply of magic states. This is an important paradigm as it is the route most commonly used in the design of fault-tolerant quantum computers.

However, stabilizer operations with access to a restricted number of magic states do not lead to universal quantum computation. Rather, the computational power depends on the type and quantity of magic states provided. It is precisely this question of computational power that we quantify by studying the complexity of simulating computations with a limited resource of magic states.

B. Magic monotone definitions

We now introduce several magic monotones of interest, borrowing some results from the general resource-theory literature. Although in our discussion we specialize to the theory of magic states, the basic considerations below can also be applied to more general resources in which the set of free states is defined by convex combinations of free pure states, which includes important examples such as coherence and entanglement. We elaborate on this in Sec. VIII.

For pure states, we define the following.

Definition 2 ([25]). The *pure-state extent* ξ is the quantity

$$\xi(\Psi) := \min\{\|c\|_1^2 : |\Psi\rangle = \sum_j c_j |\phi_j\rangle; |\phi_j\rangle \in S_n\}. \quad (7)$$

In magic theory, ξ is the stabilizer extent [25]. A related quantity appears in other resource theories such as entanglement, where it admits an analytical formula as the

squared sum of the Schmidt coefficients of a state [54], or in coherence theory, where it is the square of the ℓ_1 norm of coherence [51]. It is well known [25,55] that this can be recast as a dual optimization problem

$$\xi(\Psi) := \max \{ |\langle \omega | \Psi \rangle|^2 : |\langle \omega | \phi \rangle| \leq 1 \forall |\phi\rangle \in S_n \}. \quad (8)$$

Here, we define an ω witness to be any feasible solution to the optimization problem in Eq. (8).

We now consider four monotones, of which three can be regarded as mixed-state extensions of ξ . First, one can extend the extent to mixed states using a convex roof extension [56].

Definition 3. The *mixed-state extent* Ξ is the quantity

$$\Xi(\rho) := \min \left\{ \sum_j p_j \xi(\Psi_j) : \rho = \sum_j p_j |\Psi_j\rangle\langle\Psi_j| \right\},$$

where every $|\Psi_j\rangle$ is a pure state and p_j are non-negative coefficients such that $\sum_j p_j = 1$. Furthermore, if the minimum can be achieved with a decomposition where all $\xi(\Psi_j)$ are equal, then we say the state admits an *equimagical decomposition*.

We also consider quasiprobability distributions over free states as follows.

Definition 4 ([21]). The *robustness* \mathcal{R} is the quantity

$$\mathcal{R}(\rho) := \min \left\{ \|q\|_1 : \rho = \sum_j q_j |\phi_j\rangle\langle\phi_j|; |\phi_j\rangle \in S_n \right\},$$

where q_j are real coefficients.

In magic theory, \mathcal{R} is called the robustness of magic [21,41], inspired by the (standard) robustness of entanglement [57]. This quantity is precisely the negativity with respect to the frame defined by the set of pure-state stabilizer projectors. In particular, the robustness uses decompositions where the rank-one ket-bra terms are Hermitian. Relaxing this, we have the following definition.

Definition 5. The *dyadic negativity* Λ is the quantity

$$\Lambda(\rho) := \min \left\{ \|\alpha\|_1 : \rho = \sum_j \alpha_j |L_j\rangle\langle R_j|; |L_j\rangle, |R_j\rangle \in S_n \right\},$$

where the coefficients α_j are complex numbers.

The name reflects the fact that each $|L_j\rangle\langle R_j|$ comprises of a pair of vectors, and so is a dyad. Within the resource

theory of entanglement, a related quantity called the projective tensor norm was considered [54,58], and in the resource theory of coherence the dyadic negativity corresponds to the ℓ_1 norm of coherence [51]. Viewing this quantity as the primal solution of a convex optimization problem, it is useful to state the equivalent dual formulation [55] in terms of witness operators. We define the set of W witnesses, denoted \mathcal{W} , to be the Hermitian operators such that

$$\mathcal{W} := \{ W : |\langle L|W|R\rangle| \leq 1 \forall |L\rangle, |R\rangle \in S_n \}, \quad (9)$$

which by strong duality leads to

$$\Lambda(\rho) = \max\{\text{Tr}[W\rho] : W \in \mathcal{W}\}. \quad (10)$$

This brings us to our last monotone of interest.

Definition 6. The *generalized robustness* Λ^+ is the quantity

$$\Lambda^+(\rho) = \max\{\text{Tr}[W\rho] : W \in \mathcal{W}; W \geq 0\}, \quad (11)$$

where \mathcal{W} is the set of W witnesses.

A corresponding quantity to Λ^+ was first defined in entanglement theory [57,59] and appears in many resource theories.

Notice that this is similar to the dual formulation given in Eq. (10) except we further restrict to witnesses that are also positive semidefinite operators. We define a W^+ witness to be any feasible solution to the optimization problem in Eq. (11). Since W^+ witnesses are positive semidefinite, the condition $|\langle L|W|R\rangle| \leq 1$ simplifies to $\langle\psi|W|\psi\rangle \leq 1$ for all $|\psi\rangle \in S_n$.

Furthermore, the primal form of this monotone is

$$\Lambda^+(\rho) = \min\{\lambda : \rho \leq \lambda\sigma, \sigma \in \bar{S}_n\} \quad (12)$$

$$= \min \left\{ \lambda \geq 1 : \frac{\rho + (\lambda - 1)\rho'}{\lambda} \in \bar{S}_n \right\}, \quad (13)$$

where the optimization in the second line is over all density matrices ρ' . This form motivates the name of generalized robustness: rearranging Definition 4, the robustness \mathcal{R} can be similarly expressed as

$$\frac{\mathcal{R}(\rho) + 1}{2} = \min \left\{ \lambda \geq 1 : \frac{\rho + (\lambda - 1)\sigma}{\lambda} \in \bar{S}_n, \sigma \in \bar{S}_n \right\}, \quad (14)$$

where now the states in the optimization are restricted to free states in \bar{S}_n .

We stress that both Λ and Λ^+ are computable, in the sense that their evaluation corresponds to convex optimization problems—a second-order cone program for Λ , and a semidefinite program for Λ^+ —which can be evaluated using numerical solvers [60]. In practice, we are able to compute Λ^+ up to $n = 4$ and Λ up to $n = 3$, but one can certainly hope to make further progress in computing the quantities for states obeying some symmetry, just as in the case of \mathcal{R} [41]. The evaluation of convex roof-based quantities such as Ξ is notoriously hard in general [61], although one could again use symmetry to facilitate it in special cases [62]. Our results in Secs. III–IV further simplify the computation of all of the monotones for the practically important case of tensor products of single-qubit states.

The monotones have been considered from the perspective of general resource theories [55], and in particular they have been shown to satisfy a number of useful properties:

1. *faithfulness*: $\mathcal{M}(\rho) = 1$ if and only if $\rho \in \bar{S}_n$;
2. *monotonicity*: $\mathcal{M}(\rho) \geq \mathcal{M}[O(\rho)]$ for any free operation $O \in \mathcal{O}_n$;
3. *strong monotonicity* (monotonicity on average under selective free measurements):

$$\mathcal{M}(\rho) \geq \sum_i p_i \mathcal{M} \left(\frac{K_i \rho K_i^\dagger}{p_i} \right), \quad (15)$$

where $\{K_i\}_i$ are the Kraus operators of a quantum channel such that each K_i is stabilizer preserving, i.e., $K_i|\phi\rangle \propto |\phi'\rangle \in S_n \forall |\phi\rangle \in S_n$, and $p_i = \text{Tr}(K_i \rho K_i^\dagger)$;

4. *convexity*: $\mathcal{M}(\sum_j p_j \rho_j) \leq \sum_j p_j \mathcal{M}(\rho_j)$;
5. *submultiplicativity*: $\mathcal{M}(\otimes_j \rho_j) \leq \prod_j \mathcal{M}(\rho_j)$.

We remark that, although \mathcal{R} and Λ^+ are monotones in any convex resource theory, the fact that Λ and Ξ obey monotonicity under all completely stabilizer-preserving operations \mathcal{O}_n is a consequence of two properties: the strong monotonicity of the measures [55] coupled with the fact that any operation $O \in \mathcal{O}_n$ can be expressed in terms of Kraus operators $\{K_i\}_i$, which preserve the set of stabilizer states [22]. If we instead work with logarithmic monotones, $\mathcal{M}_{\log}(\rho) = \log[\mathcal{M}(\rho)]$ then multiplicativity becomes additivity, faithfulness instead has a $\mathcal{M}_{\log}(\rho) = 0$ condition, and due to concavity of the logarithm \mathcal{M}_{\log} is no longer a convex function but still obeys strong monotonicity [63]. Here we find it convenient to work without the logarithm in most cases.

Next, we present some general relations between these monotones that are reminiscent of known results in general resource theories [55].

Lemma 1 ([55]). *For any pure state*

$$\Lambda^+(|\Psi\rangle\langle\Psi|) = \Lambda(|\Psi\rangle\langle\Psi|) = \Xi(|\Psi\rangle\langle\Psi|) = \xi(\Psi). \quad (16)$$

Therefore, our monotones can be interpreted as mixed-state extensions of ξ . We also observe the following.

Theorem 2. *For any state ρ we have*

$$\Lambda^+(\rho) \leq \Lambda(\rho) \leq \Xi(\rho). \quad (17)$$

For completeness, we provide alternative proofs of these results in Appendix A1. Since Λ^+ is often easier to evaluate than Λ and $\Lambda^+ \leq \Lambda$, in practical settings, one can approximate Λ by evaluating Λ^+ .

C. Connecting monotones with simulation

To further motivate our investigation of the magic monotones that follows in the subsequent sections, we summarize our main results and show how the properties of the monotones will be vital to the understanding of several classes of classical simulation algorithms. Our first simulation algorithm is a quasiprobability-based approach, which introduces several novel modifications to standard Monte Carlo techniques, notably the use of dyadic frames.

Theorem 14 (informal). *Consider an n -qubit initial state with known decomposition into dyads $\rho = \sum_j \alpha_j |L_j\rangle\langle R_j|$ where $\|\alpha\|_1 = \Lambda(\rho)$. Let \mathcal{E} be a sequence of T stabilizer-preserving operations, each acting on a few qubits. Then, given a stabilizer projector Π , we can estimate the Born-rule probability $\mu = \text{Tr}(\Pi \mathcal{E}[\rho])$ with probability $1 - p_{\text{fail}}$ and additive error ϵ within a runtime*

$$\frac{\Lambda(\rho)^2}{\epsilon^2} \log(p_{\text{fail}}^{-1}) T \text{poly}(n). \quad (18)$$

Hence, the dyadic negativity Λ exactly characterizes our algorithm’s runtime. To understand how the performance scales when more copies of the input state ρ are provided, it is then necessary to understand the multiplicativity of Λ . We solve this question completely with the following.

Theorem 10. *Let σ_j be single-qubit states. Then*

$$\Lambda(\otimes_j \sigma_j) = \Xi(\otimes_j \sigma_j) = \Lambda^+(\otimes_j \sigma_j) = \prod_j \Lambda^+(\sigma_j). \quad (19)$$

This not only reveals a connection between three monotones introduced previously—allowing, for instance, for the evaluation of the generally hard-to-compute quantifier Ξ —but also shows them to be strictly multiplicative for qubit states. Consequently, when we plot these quantities on a log scale, we get a straight line, as shown with the example in Fig. 1. Although a common occurrence




FIG. 1. The scaling of magic monotones for many copies of a noisy single-qubit magic state ρ , highlighting several of our results. Our new monotones, $\Lambda^+(\rho^{\otimes n})$, $\Lambda(\rho^{\otimes n})$, and $\Xi(\rho^{\otimes n})$ are proved to be equal with multiplicative scaling leading to a straight line (gray) on this logarithmic scale (due to Theorem 10). We contrast this with a previously studied monotone, the robustness of magic \mathcal{R} , for which we can numerically compute the value up to $n = 5$ (shown as purple data points). For $n > 5$, the shaded purple region shows the possible values of \mathcal{R} as enforced by upper bounds (due to sub-multiplicativity) and two lower bounds (Lemma 11 and Theorem 12). The robustness of magic has a wide range of possible values, but even the lower bound grows exponentially faster than the value of our new monotones, entailing that classical simulation algorithms based on the new monotones offer an improvement in the exponential scaling of their runtime.

in the structurally simpler theory of qudit magic states [35,64], multiplicativity has not been shown before for any mixed-state monotone in qubit magic theory.

Theorem 10 lets us avoid the main problem, which hinders an understanding of the performance of previous quasiprobability simulation algorithms such as the Howard-Campbell simulator based on the robustness \mathcal{R} , namely the inability to efficiently compute $\mathcal{R}(\rho^{\otimes n})$ for large n [21,41]. In addition, we can use the multiplicativity result to show an exponential separation between our monotones and the robustness of magic.

Theorem 12. *Given any single-qubit nonstabilizer state ρ , there exists positive real constants α and β where $\alpha > \beta$ and so that*

$$2^{\alpha n} \leq \mathcal{R}(\rho^{\otimes n}), \tag{20}$$

$$2^{\beta n} = \Lambda(\rho^{\otimes n}) = \Lambda^+(\rho^{\otimes n}) = \Xi(\rho^{\otimes n}). \tag{21}$$

This establishes the simulation algorithm of Theorem 14 as polynomially faster than previous quasiprobability simulators, as illustrated by the example in Fig. 1.

Our second simulation algorithm is based on the stabilizer rank, which allows it to be used for both Born-rule probability estimation and for approximately sampling

from the output distribution of a quantum circuit. Importantly, existing stabilizer-rank simulation algorithms only applied to pure states [23,25]. We extend this to mixed states through the monotone Ξ as follows.

Theorem 19 (informal). *Let ρ be a state with known mixed-state extent decomposition. Then there is a classical algorithm that approximately samples from the probability distribution associated with a sequence of Pauli measurements on ρ . Our samples come from a distribution that is δ -close in ℓ_1 norm to the actual distribution, and each sample has an expected runtime*

$$\mathbb{E}(T) = \mathcal{O}[\Xi(\rho)/\delta^3] \tag{22}$$

as long as δ is not too small. Furthermore, if ρ is a product of single-qubit states, there is no variance in the runtime.

There are two notable technical advances here: one is a factor $1/\delta$ improvement in runtime over previous simulators of this type [25], even when applied to pure states; the other improvement is the rather surprising result that sampling can often be performed without any variance in the runtime.

The last of our simulation algorithms is the constrained path simulator, which enjoys an efficient runtime, but instead sacrifices the accuracy of the simulation depending on how resourceful the input state is. The precision has an inverse polynomial dependence on the generalized robustness Λ^+ , again directly connecting a magic monotone with classical simulation.

We thus see that the tightness of our simulators' runtimes and our ability to sharply characterize them is inherited from the properties and characterization of the monotones introduced earlier. We overview the connections between the monotones and our simulation algorithms in Table I. The detailed derivation of the theorems, as well as additional results—including connecting the monotones with magic distillation rates—all follow in the remainder of the paper.

III. SINGLE-QUBIT MAGIC STATES

In this section, we present a complete description of our magic monotones for single-qubit states. Recalling from Theorem 2 that the monotones in general obey the relation $\Xi(\rho) \geq \Lambda(\rho) \geq \Lambda^+(\rho)$, the key question is then whether the inequalities can be tight, thus unifying the different approaches to the quantification of magic. We answer this in the affirmative.

Theorem 3. *For any single-qubit state ρ , we have*

$$\Lambda^+(\rho) = \Lambda(\rho) = \Xi(\rho), \tag{23}$$

and furthermore ρ admits an equimagical decomposition (recall Definition 3).

We see in the following section that this equivalence persists for tensor products of single-qubit states. However, equality does not extend to general n -qubit states for $n \geq 2$, as numerically we find that $\Lambda^+(\rho) < \Lambda(\rho)$ for most random two-qubit density matrices ρ . The proof of Theorem 3 rests on a trio of lemmata. First, we have the following.

Lemma 4 (The monotone equality lemma). *For any ω witness $|\omega\rangle$, we define the set B_ω to be the convex hull of all pure states Ψ for which $|\langle\omega|\Psi\rangle|^2 = \xi(\Psi)$. It follows that for all $\rho \in B_\omega$ we have*

$$\Lambda^+(\rho) = \Lambda(\rho) = \Xi(\rho) = \langle\omega|\rho|\omega\rangle. \quad (24)$$

Proof of Lemma 4. If $\rho \in B_\omega$, we can find a convex decomposition

$$\rho = \sum_j p_j |\Psi_j\rangle\langle\Psi_j|, \quad (25)$$

where $|\langle\omega|\Psi_j\rangle|^2 = \xi(\Psi_j)$ for all j . We can use this decomposition to obtain an upper bound on the mixed-state extent as follows:

$$\begin{aligned} \Xi(\rho) &\leq \sum_j p_j \xi(\Psi_j) \\ &= \sum_j p_j \langle\omega|\Psi_j\rangle\langle\Psi_j|\omega\rangle \\ &= \langle\omega|\rho|\omega\rangle. \end{aligned} \quad (26)$$

On the other hand, $W = |\omega\rangle\langle\omega| \in \mathcal{W}^+$ and so can be used to lower bound the generalized robustness to show

$$\langle\omega|\rho|\omega\rangle \leq \Lambda^+(\rho). \quad (27)$$

Combining Eqs. (26) and (27) with Theorem 2, we have

$$\langle\omega|\rho|\omega\rangle \leq \Lambda^+(\rho) \leq \Lambda(\rho) \leq \Xi(\rho) \leq \langle\omega|\rho|\omega\rangle. \quad (28)$$

Therefore, these inequalities all collapse to equalities. ■

Making use of Lemma 4 requires us to first understand the structure of optimal ω witnesses, which we discuss soon. However, first it is useful to define some different subsets of the Bloch sphere.

Definition 7. *The positive octant is the set*

$$P := \{\rho : \langle X \rangle, \langle Y \rangle, \langle Z \rangle \geq 0\}. \quad (29)$$

We further subdivide the positive octant as follows:




FIG. 2. The region P_Y , which is a third of the positive octant. The dotted lines show the pure states at boundaries $P_Y \cap P_X$, and $P_Y \cap P_Z$. For these boundary states, we know (by Lemma 6) that the optimal ω witness is given by Eq. (32) with the parameter set to $q = \sqrt{2/3}$. For other pure states in P_Y , the ω witness still has the form given by Eq. (32) but the parameter q may be greater than $\sqrt{2/3}$. However, interestingly, the majority of pure states in P_Y have an optimal ω witness with $q = \sqrt{2/3}$ and these are shown in yellow in this plot. On the geodesic through $|0\rangle$, $|H\rangle$, and $|+\rangle$, we have that $q = 1$. Between this geodesic and the yellow region, q varies continuously from 1 to $\sqrt{2/3}$ and this intermediate region is shown in green.

$$\begin{aligned} P_X &:= \{\rho : \rho \in P, \langle X \rangle \leq \langle Y \rangle, \langle X \rangle \leq \langle Z \rangle\}, \\ P_Y &:= \{\rho : \rho \in P, \langle Y \rangle \leq \langle X \rangle, \langle Y \rangle \leq \langle Z \rangle\}, \\ P_Z &:= \{\rho : \rho \in P, \langle Z \rangle \leq \langle X \rangle, \langle Z \rangle \leq \langle Y \rangle\}, \end{aligned} \quad (30)$$

where we use the shorthand $\langle M \rangle := \text{Tr}[\rho M]$. See Fig. 2 for an illustration of P_Y .

The sets P_X , P_Y , and P_Z further divide the positive octant into thirds and it is easy to verify that $P = P_X \cup P_Y \cup P_Z$. These sets are not quite disjoint because of the following proposition.

Proposition 5. *From Definition 7, we have the following:*

$$\begin{aligned} P_Z \cap P_Y &= \{\rho : \rho \in P, \langle Z \rangle = \langle Y \rangle \leq \langle X \rangle\}, \\ P_X \cap P_Z &= \{\rho : \rho \in P, \langle X \rangle = \langle Z \rangle \leq \langle Y \rangle\}, \\ P_X \cap P_Y &= \{\rho : \rho \in P, \langle X \rangle = \langle Y \rangle \leq \langle Z \rangle\}. \end{aligned}$$

This is straightforward to prove. For example, in P_Z the smallest expectation value is for Z and for P_Y the smallest expectation value is for Y . Therefore, in the intersection these two expectation values must be equal. We note that any state is Clifford equivalent to a state in the positive

octant P . Furthermore, the Clifford

$$F := \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & -i \\ 1 & i \end{pmatrix}, \quad (31)$$

satisfies $FXF^\dagger = Y$, $FYF^\dagger = Z$, and $FZF^\dagger = X$. Therefore, the sets P_X , P_Y , and P_Z are Clifford equivalent and therefore every state is Clifford equivalent to some $\rho \in P_Y$.

Now we are ready to characterize optimal ω witnesses.

Lemma 6. *Let $|\Psi\rangle$ be any pure, single-qubit nonstabilizer state in the set P_Y . Then the ω witness $|\omega\rangle$ that achieves $|\langle\Psi|\omega\rangle|^2 = \xi(\Psi)$ has an operator representation of the form*

$$|\omega\rangle\langle\omega| = \frac{\mathbb{1} + qH + \sqrt{1 - q^2}Y}{1 + q/\sqrt{2}}, \quad (32)$$

where $\sqrt{2/3} \leq q \leq 1$ and $H = (X + Z)/\sqrt{2}$. Furthermore, if $|\Psi\rangle$ is in the set $P_Y \cap P_X$ or $P_Y \cap P_Z$ then $q = \sqrt{2/3}$ and the ω witness takes the form

$$|\omega\rangle\langle\omega| = \frac{\mathbb{1} + (X + Y + Z)/\sqrt{3}}{1 + 1/\sqrt{3}}. \quad (33)$$

The actual value of the variable q is easy to numerically compute, but is analytically complicated and not instructive to present. Rather, in Fig. 2, we illustrate the region P_Y and highlight where $q = \sqrt{2/3}$ and $q > \sqrt{2/3}$.

Proof of Lemma 6. We begin by observing that for any $|\Psi\rangle$ there exists a decomposition into stabilizer states such that $|\Psi\rangle = \sum_j c_j |\phi_j\rangle$ and $\xi(\Psi) = \left(\sum_j |c_j|\right)^2$. Given an optimal ω witness we have

$$\xi(\Psi) = |\langle\omega|\Psi\rangle|^2 = \left| \sum_j c_j \langle\omega|\phi_j\rangle \right|^2. \quad (34)$$

Therefore,

$$\left(\sum_j |c_j| \right)^2 = \left| \sum_j c_j \langle\omega|\phi_j\rangle \right|^2. \quad (35)$$

Given that $|\langle\omega|\phi_j\rangle| \leq 1$, the above equality can only hold if $|\langle\omega|\phi_j\rangle| = 1$ for every j with $|c_j| > 0$. In particular, if Ψ is a nonstabilizer state it must have at least two nonzero c_j terms, and there must exist at least two stabilizer states such that $|\langle\omega|\phi_j\rangle| = 1$. We return to use this fact shortly.

Using the set of Pauli matrices as a basis

$$|\omega\rangle\langle\omega| = \lambda(\mathbb{1} + q_x X + q_y Y + q_z Z), \quad (36)$$

where the coefficients q_x, q_y , and q_z are real. Since $|\omega\rangle\langle\omega|$ is a rank-1 operator we know

$$q_x^2 + q_y^2 + q_z^2 = 1, \quad (37)$$

and since $|\omega\rangle\langle\omega|$ is a positive operator we have $\lambda > 0$. Given a valid ω witness, we can always obtain another valid ω witness by permuting any of $\{q_x, q_y, q_z\}$ or changing the signs. Therefore, the optimal ω witness for a state in the set P_Y has $q_x \geq q_y$ and $q_z \geq q_y$, since this ordering maximizes $|\langle\Psi|\omega\rangle|^2$. This means that the two stabilizer states with the largest overlap with $|\Psi\rangle$ are $|+\rangle$ and $|0\rangle$. We show earlier there must be at least two stabilizer states for which $|\langle\omega|\phi_j\rangle| = 1$, so we conclude $|\langle\omega|+\rangle| = 1$ and $|\langle\omega|0\rangle| = 1$. It follows that $q_x = q_z$ and we define $q := \sqrt{2}q_x = \sqrt{2}q_z$. Condition Eq. (37) implies that $q_y = \sqrt{1 - q^2}$ so we show that the optimal ω witness has the form

$$|\omega\rangle\langle\omega| = \lambda \left(\mathbb{1} + q \frac{X + Z}{\sqrt{2}} + \sqrt{1 - q^2} Y \right). \quad (38)$$

Furthermore, $|\langle\omega|+\rangle| = 1$ implies that

$$\lambda = (1 + q/\sqrt{2})^{-1}. \quad (39)$$

Lastly, we note that the condition $|\langle\omega|+i\rangle| \leq 1$ entails that $q \geq \sqrt{2/3}$. Therefore, we know the form of the ω witness in the set P_Y and prove that $\sqrt{2/3} \leq q \leq 1$. Next, consider the special case when the state is at an intersection, such as $P_Y \cap P_Z$. Then, the optimal ω witness has the above form determined for the region P_Y . However, the region P_Z only differs by an F rotation, so the optimal ω witness must have a similar form but with the Pauli operators permuted, so that

$$|\omega\rangle\langle\omega| = \frac{\left(\mathbb{1} + p \frac{Y+X}{\sqrt{2}} + \sqrt{1 - p^2} Z \right)}{2(1 + p/\sqrt{2})}. \quad (40)$$

The only way Eqs. (32) and (40) can both be true, is if $q = p = \sqrt{2/3}$. A similar argument holds for $P_Y \cap P_X$ and this proves Lemma 6. ■

Our third lemma shows that every mixed state is contained in an appropriate convex set.

Lemma 7. *For any single-qubit nonstabilizer state ρ , there exists a ω witness $|\omega\rangle$ such that $\rho \in B_\omega$ (as defined in Lemma 4).*

This implies that for single-qubit states we can leverage Lemmas 4 and 6 to prove Theorem 3.

Proof of Lemma 7. We consider individual slices of the Bloch sphere such that $\text{Tr}[\rho\sigma_F] = f$ where $\sigma_F = (X + Y + Z)/\sqrt{3}$ and f is equal to the inner product between the Bloch vectors representing ρ and σ_F . A particular f value specifies a slice through the Bloch sphere. Let us denote \mathbb{S}_f as the set of all states inside this slice. For every nonstabilizer state in the positive octant we have $1/\sqrt{3} \leq f$, and for all normalized states we have $f \leq 1$. Within this slice there are three special, pure states, which are

$$\begin{aligned} |\Psi_f^X\rangle\langle\Psi_f^X| &:= \frac{1}{2}(\mathbb{1} + aY + aZ + \sqrt{1 - 2a^2}X), \\ |\Psi_f^Y\rangle\langle\Psi_f^Y| &:= \frac{1}{2}(\mathbb{1} + aX + aZ + \sqrt{1 - 2a^2}Y), \\ |\Psi_f^Z\rangle\langle\Psi_f^Z| &:= \frac{1}{2}(\mathbb{1} + aY + aX + \sqrt{1 - 2a^2}Z), \end{aligned} \quad (41)$$

where a obeys

$$\sqrt{3}f = 2a + \sqrt{1 - 2a^2}. \quad (42)$$

For $1/\sqrt{3} \leq f \leq 1$, there is a unique a such that $a \in [0, 1/\sqrt{3}]$ and $\sqrt{1 - 2a^2} \in [1/\sqrt{3}, 1]$. Crucially, these states are the unique pure states of the following set intersections.

$$\begin{aligned} |\Psi_f^X\rangle\langle\Psi_f^X| &\in \mathbb{S}_f \cap P_Y \cap P_Z, \\ |\Psi_f^Y\rangle\langle\Psi_f^Y| &\in \mathbb{S}_f \cap P_X \cap P_Z, \\ |\Psi_f^Z\rangle\langle\Psi_f^Z| &\in \mathbb{S}_f \cap P_X \cap P_Y. \end{aligned} \quad (43)$$

Referring back to Proposition 5, it is clear that these states must have the form given in Eq. (41).

Notice that these special states are Clifford rotations of each other. By Lemma 6 these three special states all have the same optimal ω witness given by Eq. (33). Since they share their optimal ω witness, Lemma 4 applies to all convex combinations of states $\{\Psi_f^X, \Psi_f^Y, \Psi_f^Z\}$ as illustrated in Fig. 3. Note that $\{\Psi_f^X, \Psi_f^Y, \Psi_f^Z\}$ all have the same value for the extent, since

$$\xi(\Psi_f^X) = \xi(\Psi_f^Y) = \xi(\Psi_f^Z) = \frac{1 + f}{1 + 1/\sqrt{3}}. \quad (44)$$

Therefore, for these states, a mixture of states with the same amount of magic achieves the optimal convex-roof extension. That is, each of these states admit an equimagic decompositions.

Next, we consider mixed states outside the convex hull of $\{\Psi_f^X, \Psi_f^Y, \Psi_f^Z\}$ and inside $P_Y \cap \mathbb{S}_f$ as illustrated in




FIG. 3. A slice \mathbb{S}_f through the positive octant. States are parameterized by the coordinates $\{r_A, r_B, r_F\}$ as defined in Eq. (46). For the slice \mathbb{S}_f we have $r_F = f$ for some constant f . The axes for the $\{r_A, r_B\}$ coordinates are shown in red. The slice is divided into thirds corresponding to the sets defined in Definition 7 with the intersections of these sets shown with dashed lines. Where these intersections meet the pure states we label the pure states $\{\Psi_f^X, \Psi_f^Y, \Psi_f^Z\}$ defined in Eq. (43), and the purple triangle denotes the convex hull of the set $\{\Psi_f^X, \Psi_f^Y, \Psi_f^Z\}$. States outside this convex set are considered as a mixture of two pure states $|\Phi_\rho^\pm\rangle$, defined in Eq. (49) and shown with green dots.

Fig. 3. We define a set of linearly independent, Hermitian operators

$$\sigma_A = \frac{X + Z - 2Y}{\sqrt{6}}, \quad \sigma_B = \frac{X - Z}{\sqrt{2}}, \quad (45)$$

and σ_F as defined earlier. The set $\{\sigma_A, \sigma_B, \sigma_F\}$ is unitarily equivalent to $\{X, Y, Z\}$, so every state can be decomposed as

$$\rho = (\mathbb{1} + r_A\sigma_A + r_B\sigma_B + r_F\sigma_F)/2, \quad (46)$$

where inside the slice \mathbb{S}_f we have $r_F = f$. The variables r_A and r_B are used for the coordinate system in Fig. 3.

Given a mixed state ρ , we can define a pair of pure states Φ_ρ^+ and Φ_ρ^- , such that

$$\begin{aligned} \langle\Phi_\rho^\pm|\sigma_A|\Phi_\rho^\pm\rangle &= \text{Tr}[\sigma_A\rho] = r_A, \\ \langle\Phi_\rho^\pm|\sigma_F|\Phi_\rho^\pm\rangle &= \text{Tr}[\sigma_F\rho] = f, \end{aligned} \quad (47)$$

and the states are pure, so that

$$\langle\Phi_\rho^\pm|\sigma_A|\Phi_\rho^\pm\rangle^2 + \langle\Phi_\rho^\pm|\sigma_B|\Phi_\rho^\pm\rangle^2 + \langle\Phi_\rho^\pm|\sigma_F|\Phi_\rho^\pm\rangle^2 = 1. \quad (48)$$

There are two possible solutions for $\langle\Phi_\rho^\pm|\sigma_B|\Phi_\rho^\pm\rangle$, which leads to

$$|\Phi_\rho^\pm\rangle\langle\Phi_\rho^\pm| = \frac{1}{2} \left(\mathbb{1} + r_A\sigma_A \pm \sqrt{1 - r_A^2 - f^2}\sigma_B + f\sigma_F \right). \quad (49)$$

By construction, ρ is a convex combination of Φ_ρ^+ and Φ_ρ^- . The geometry is illustrated in Fig. 3, where the pair of purified states are shown as green dots with ρ located on the line between them. To deploy Lemma 4, it remains to prove that Φ_ρ^\pm share an optimal ω witness.

The states Φ_ρ^\pm are both in the region P_Y , which can be seen from the geometry on Fig. 3 though we also give an algebraic proof in Appendix B. Due to $\Phi_\rho^\pm \in P_Y$, we can use Lemma 6 to determine the form of their optimal ω witnesses. In Lemma 6, the witness $\omega(q)$ had a free parameter q that we had to maximize over. Since $\langle \omega(q) | \sigma_B | \omega(q) \rangle = 0$ for any q value, we have

$$\langle \omega(q) | \rho | \omega(q) \rangle = |\langle \Phi_\rho^+ | \omega(q) \rangle|^2 = |\langle \Phi_\rho^- | \omega(q) \rangle|^2. \quad (50)$$

Performing the maximization over q , the optimal q value is the same for Φ_ρ^+ and Φ_ρ^- due to Eq. (50). Therefore, Φ_ρ^\pm share exactly the same optimal ω witness. This completes the proof of Lemma 7 and thus also of Theorem 3. ■

IV. MULTIPLICATIVITY

We now study the behavior of the monotones Ξ , Λ , and Λ^+ for tensor products of states. It was found by Bravyi *et al.* [25] that ω witnesses of small dimension are closed under tensor products, formalized as follows.

Theorem 8 ([25]). *Let $|\omega_j\rangle$ be vectors from a one-, two-, or three-qubit Hilbert space such that each ω_j is an ω witness. Then $|\Omega\rangle := \otimes_j |\omega_j\rangle$ is an ω witness.*

This is a rewording of Corollary 1 and Corollary 3 of Ref. [25]. From the above result, Ref. [25] further showed that the extent is multiplicative for such tensor products.

Theorem 9 ([25]). *Let $|\psi_j\rangle$ be one-, two-, or three-qubit states. Then*

$$\xi(\otimes_j |\psi_j\rangle) = \prod_j \xi(|\psi_j\rangle). \quad (51)$$

Here, we give a related multiplicativity result for several mixed-state monotones.

Theorem 10. *Let σ_j be single-qubit states. Then*

$$\Lambda(\otimes_j \sigma_j) = \Xi(\otimes_j \sigma_j) = \Lambda^+(\otimes_j \sigma_j) = \prod_j \Lambda^+(\sigma_j), \quad (52)$$

and furthermore $\otimes_j \sigma_j$ admits an equimagical decomposition (recall Definition 3).

Prior to this work, there were no known strict multiplicativity results for resource monotones for mixed states in qubit magic theory. For instance, Howard and Campbell

[21] found that the robustness of magic can be strictly submultiplicative, $\mathcal{R}(\rho \otimes \rho) < \mathcal{R}(\rho)^2$ for all nonstabilizer ρ considered, and we discuss this later in this section. There does exist a multiplicative lower bound on the robustness of magic, proved using the so-called stabilizer norm [21]. However, the lower bounds and upper bounds appear to always be loose and so we have no strict multiplicativity results. Additionally, Raussendorf *et al.* [40] introduced a qubit-based phase-space robustness \mathcal{R}_{PS} that can behave strictly supermultiplicatively, so that $\mathcal{R}_{\text{PS}}(\rho \otimes \rho) > \mathcal{R}_{\text{PS}}(\rho)^2$ for some ρ . It is natural to wonder if Theorem 9 or Theorem 10 could extend to a tensor product of states with arbitrary dimension. However, in the final stages of completing this work, it was proved that this cannot hold in full generality [46]. It remains an open question whether the monotones satisfy multiplicativity for states composed of a low number of qubits, mirroring the multiplicativity of the extent; indeed, numerical results suggest that Λ^+ is also multiplicative for mixed two-qubit states.

Proof of Theorem 10. From the definition of Ξ we see that it is manifestly submultiplicative. Combining this observation with Theorem 3 we have that

$$\Xi(\otimes_j \sigma_j) \leq \prod_j \Xi(\sigma_j) = \prod_j \Lambda^+(\sigma_j) \quad (53)$$

holds for all products of single-qubit states. Strengthening this to strict equality requires us to find a matching lower bound. The proof of Theorem 3 establishes that for every single-qubit state the optimal W^+ witness has the form $|\omega_j\rangle\langle\omega_j|$ where ω_j is an ω witness. By Theorem 8, $|\Omega\rangle = \otimes_j |\omega_j\rangle$ is also an ω witness, and consequently $|\Omega\rangle\langle\Omega| = \otimes_j |\omega_j\rangle\langle\omega_j|$ is a W^+ witness that can be used to lower bound Λ^+ as follows:

$$\prod_j \Lambda^+(\sigma_j) = \langle \Omega | \otimes_j \sigma_j | \Omega \rangle \leq \Lambda^+(\otimes_j \sigma_j). \quad (54)$$

Combining Eq. (53), Eq. (54), and Theorem 2 we obtain

$$\begin{aligned} \prod_j \Lambda^+(\sigma_j) &\leq \Lambda^+(\otimes_j \sigma_j) \\ &\leq \Lambda(\otimes_j \sigma_j) \leq \Xi(\otimes_j \sigma_j) \leq \prod_j \Lambda^+(\sigma_j). \end{aligned} \quad (55)$$

Since the left- and rightmost quantities are the same, all these inequalities must collapse to equalities.

It remains to show that these product states admit equimagical decompositions. This is easily verified by taking an equimagical decomposition for each single qubit state (existence ensured by Theorem 3) and using this to construct the natural decomposition for the product state. It then follows immediately from Theorem 9 that each pure term has equal extent and by the above argument that this is optimal with respect to the Ξ monotone. ■

V. COMPARISON WITH ROBUSTNESS

Here we discuss how our new monotones scale compared to the robustness of magic (recall Definition 4). While Λ^+ , Λ , and Ξ are often equal, the robustness of magic is typically much larger, as formalized in the following result.

Lemma 11. *For any density matrix ρ we have*

$$\mathcal{R}(\rho) \geq 2\Lambda^+(\rho) - 1. \quad (56)$$

Furthermore, if ρ is a single-qubit state this tightens to

$$\mathcal{R}(\rho) \geq (1 + \sqrt{2})\Lambda^+(\rho) - \sqrt{2}. \quad (57)$$

We remark that a similar result to Eq. (56) for Λ is claimed in Refs. [55,58], but the proof contains an error.

However, because the robustness of magic is not multiplicative, Lemma 11 does not tell us much about how the different monotones scale. For this, we observe that the gap can scale exponentially.

Theorem 12. *Given any single-qubit nonstabilizer state ρ , there exists positive real constants α and β where $\alpha > \beta$ and*

$$2^{\alpha n} \leq \mathcal{R}(\rho^{\otimes n}), \quad (58)$$

$$2^{\beta n} = \Lambda(\rho^{\otimes n}) = \Lambda^+(\rho^{\otimes n}) = \Xi(\rho^{\otimes n}). \quad (59)$$

For example, for the Hadamard $|H\rangle$ state we show that this holds with $\alpha = 0.271\,553$ and $\beta = 0.228\,443$.

Proof of Lemma 11. The dual formulation of the robustness of magic tells us that $\mathcal{R}(\rho) \geq \text{Tr}[R\rho]$ for any R such that $|\langle\phi|R|\phi\rangle| \leq 1$ for all $|\phi\rangle$ that are stabilizer states. We call such an operator an R witness. Note that an R witness is not necessarily positive. Let W denote the W^+ witness such that $\Lambda^+(\rho) = \text{Tr}[W\rho]$. Now, we consider the operator

$$R = \frac{2}{1-s}W - \frac{1+s}{1-s}\mathbb{1}, \quad (60)$$

where

$$s = \min_{\phi \in \mathcal{S}_n} \langle\phi|W|\phi\rangle. \quad (61)$$

Next, we show R is indeed an R witness. For any $|\phi\rangle \in \mathcal{S}_n$,

$$\begin{aligned} \langle\phi|R|\phi\rangle &= \frac{2}{1-s}\langle\phi|W|\phi\rangle - \frac{1+s}{1-s} \\ &\leq \frac{2}{1-s} - \frac{1+s}{1-s} = 1, \end{aligned} \quad (62)$$

where we use $\langle\phi|W|\phi\rangle \leq 1$. Using $\langle\phi|W|\phi\rangle \geq s$, we similarly obtain

$$\langle\phi|R|\phi\rangle \geq \frac{2s}{1-s} - \frac{1+s}{1-s} = -1. \quad (63)$$

Therefore, R is indeed an R witness and we can lower bound the robustness as follows:

$$\begin{aligned} \mathcal{R}(\rho) \geq \text{Tr}[R\rho] &= \frac{2}{1-s}\Lambda^+(\rho) - \frac{1+s}{1-s} \\ &= \frac{[2\Lambda^+(\rho) - 1] - s}{1-s}, \end{aligned} \quad (64)$$

Since $\Lambda^+(\rho) \geq 1$, the right-hand side is monotonically increasing with s on the relevant range $s \in [0, 1)$. This prompts the question whether we can lower bound s . By definition $s \geq 0$ for any W^+ witness and so Eq. (56) holds in general. In the special case of single-qubit states, and assuming for brevity that $\rho \in P_Y$, we know the optimal witness has the form $W = |\omega\rangle\langle\omega|$ given by Lemma 6. Since $|0\rangle$ has the largest possible overlap with $|\omega\rangle$, it follows that $|1\rangle$ must have the smallest possible overlap and one finds that

$$\langle 1|W|1\rangle = s = (1 - q/\sqrt{2})/(1 + q/\sqrt{2}). \quad (65)$$

Over the allowed range $q \in [\sqrt{2/3}, 1]$, we have

$$s \geq (1 - 1/\sqrt{2})/(1 + 1/\sqrt{2}), \quad (66)$$

for every optimal single-qubit W^+ witness. Substituting this into Eq. (64) gives Eq. (57). ■

Proof of Theorem 12. The stab norm \mathcal{D} has been shown to provide a lower bound on the robustness of magic (see the Supplemental Material of Ref. [21] and also Ref. [65]), so that for any single-qubit nonstabilizer state ρ we have

$$\mathcal{D}(\rho)^n = \mathcal{D}(\rho^{\otimes n}) \leq \frac{\mathcal{D}(\rho)^n - \frac{1}{2^n}}{1 - \frac{1}{2^n}} \leq \mathcal{R}(\rho^{\otimes n}), \quad (67)$$

where the stab norm of a single-qubit state is

$$\mathcal{D}(\rho) = \frac{1}{2}(1 + |\langle X \rangle| + |\langle Y \rangle| + |\langle Z \rangle|). \quad (68)$$

Defining $\alpha = \log_2(\mathcal{D}(\rho))$, we obtain Eq. (58). For instance $\mathcal{D}(|H\rangle\langle H|) = 1.207$ and so $\alpha = 0.271\,553$ for Hadamard states.

Similarly, Eq. (59) holds due to Theorem 10 and setting $\beta := \log_2[\Lambda(\rho)]$. For instance, $\beta = 0.228\,443$ for Hadamard states. To show $\alpha > \beta$ in general, we need to

show that $\mathcal{D}(\rho) > \Lambda(\rho)$ for all nonstabilizer, single-qubit states. We note that for a single-qubit we have

$$\mathcal{R}(\rho) = |\langle X \rangle| + |\langle Y \rangle| + |\langle Z \rangle| \quad (69)$$

for any nonstabilizer state. This can be shown by using Eq. (67) to obtain a lower bound on $\mathcal{R}(\rho)$, with the corresponding upper bound following from a simple quasiprobability decomposition into stabilizer states. Therefore, $\mathcal{R}(\rho) = 2\mathcal{D}(\rho) - 1$ and combining this with Lemma 11, we get

$$\mathcal{D}(\rho) \geq \frac{1 + \sqrt{2}}{2} \Lambda^+(\rho) - \left(\frac{\sqrt{2} - 1}{2} \right). \quad (70)$$

This reveals that $\mathcal{D}(\rho) > \Lambda^+(\rho)$ whenever $\Lambda^+(\rho) > 1$. ■

We further remark that the robustness of magic is not multiplicative and the known upper bounds on $\mathcal{R}(\rho^{\otimes n})$ are loose compared to the lower bound in Eq. (58). For instance, Heinrich and Gross [41] showed that for the Hadamard state (or the equivalent T state) $\mathcal{R}(|H\rangle\langle H|^{\otimes n}) = \mathcal{O}(2^{0.368601n})$ and this is the best known upper bound.

VI. DISTILLATION AND ASYMPTOTIC RATES

We now consider the scenario of distillation—that is, consuming many copies of an input resource state ρ to prepare copies of some target state—and show how the quantifiers we introduce characterize this task. Firstly, it is easy to see using the multiplicativity of the magic monotones Λ , Λ^+ , and Ξ for single-qubit systems together with their monotonicity that, whenever there exists a stabilizer operation taking $\rho^{\otimes k} \rightarrow \sigma^{\otimes m}$ for some single-qubit ρ and σ , we must have

$$\frac{k}{m} \geq \frac{\log \Lambda(\sigma)}{\log \Lambda(\rho)}, \quad (71)$$

and analogously for the other magic monotones. This already allows one to obtain insightful no-go results on the transformations between stabilizer states and gate synthesis, along the lines considered in Ref. [21] but without the need to perform the difficult computation of the monotones for many copies of a state.

However, in practical settings it is often desirable to go beyond such exact transformations and consider protocols that allow for imperfect conversion. Our quantifiers can yield bounds for the efficiency of more general distillation protocols and their asymptotic rates. We focus on the magic monotone Λ^+ as it is the most efficiently computable out of the three and gives us the tightest bounds. A useful property of Λ^+ is its monotonicity on average under

general probabilistic protocols: specifically, we have [55]

$$\Lambda^+(\rho) \geq \sum_i p_i \Lambda^+ \left[\frac{O_i(\rho)}{p_i} \right], \quad (72)$$

where each O_i is a stabilizer-preserving quantum operation that need not preserve trace [i.e., $O_i(\sigma) \propto \omega \in \mathcal{S}_n \forall \sigma \in \bar{\mathcal{S}}_n$], the overall quantum operation $\sum_i O_i$ preserves trace, and $p_i = \text{Tr}[O_i(\rho)]$ denotes the probability that the input state ρ is transformed to the output $O_i(\rho)$.

The most general representation of a distillation protocol is then an operation, which takes k copies of a given input ρ to m copies of some desired pure output state ψ , up to error ε in fidelity, and succeeding with probability p . All such protocols are limited as follows.

Theorem 13. *Let ρ be any n -qubit quantum state, and ψ a pure state of at most three qubits. If there exists a probabilistic (that is, not necessarily trace-preserving) stabilizer operation taking $\rho^{\otimes k} \rightarrow p\tau$, where τ is a state such that $\langle \psi^{\otimes m} | \tau | \psi^{\otimes m} \rangle \geq 1 - \varepsilon$, then it necessarily holds that*

$$k \geq \frac{\log p + \log(1 - \varepsilon) + m \log F(\psi)^{-1}}{\log \Lambda^+(\rho)} \quad (73)$$

and

$$k \geq p \left[\frac{\log(1 - \varepsilon) + m \log F(\psi)^{-1}}{\log \Lambda^+(\rho)} \right], \quad (74)$$

where $F(\psi) = \max_{|\phi\rangle \in \mathcal{S}_n} |\langle \psi | \phi \rangle|^2$ denotes the stabilizer fidelity [25].

The above establishes two bounds on the least number of copies of ρ necessary to perform the distillation of ψ up to the desired accuracy, characterizing the dependence on the resources contained in both ρ (as quantified by Λ^+) and in ψ (as quantified by stabilizer fidelity F). Note that either of the two bounds can perform better, depending on the values of the parameters (see Fig. 4).

Proof. By submultiplicativity of Λ^+ we have $\Lambda^+(\rho)^k \geq \Lambda^+(\rho^{\otimes k})$. By monotonicity black under probabilistic protocols [see Eq. (72)] we have $\Lambda^+(\rho^{\otimes k}) \geq p \Lambda^+(\tau)$. Because $|\psi\rangle\langle\psi|^{\otimes m} / F(\psi^{\otimes m})$ is a W^+ witness and hence a feasible solution to the dual form of Λ^+ , we arrive at

$$\begin{aligned} \Lambda^+(\rho)^k &\geq p \text{Tr} \left[\tau \frac{|\psi\rangle\langle\psi|^{\otimes m}}{F(\psi^{\otimes m})} \right] \\ &\geq p \frac{1 - \varepsilon}{F(\psi^{\otimes m})}. \end{aligned} \quad (75)$$

If ψ is any single-qubit, two-qubit, or three-qubit pure state, then $F(\psi^{\otimes m}) = F(\psi)^m$ (see Ref. [25] or Theorem




FIG. 4. Comparison of the lower bounds for the number of copies k of the state $\rho = \alpha|H\rangle\langle H| + (1 - \alpha)\mathbb{1}/2$ necessary to distill m copies of $|H\rangle$ with success probability $p = 0.9$ and output infidelity ε . In (a), we fix $\alpha = 0.75$ and demonstrate that the bounds in this paper can characterize distillation well in a range of physical error regimes even for a small number of target copies ($m = 4$), providing a better bound than Ref. [47] down to $\varepsilon \approx 10^{-21}$. In (b), we show that the bounds substantially improve when m increases. This suggests in particular that, even though the bound of Ref. [47] gets increasingly better as $\varepsilon \rightarrow 0$ by construction, in practical regimes its performance can be exceeded by considering a larger number of copies of the distillation target. We show this in (c) by varying the input error parameter α with a fixed small output error of $\varepsilon = 10^{-20}$ and 24 target copies m . Our bounds perform better even in the regime of α close to 1, and their performance can be improved further by considering larger m . Note that our bounds apply also to the pure-state case ($\alpha = 1$), while the bound of Ref. [47] explicitly applies only to full-rank inputs.

8), and so

$$\Lambda^+(\rho)^k \geq p \frac{1 - \varepsilon}{F(\psi^{\otimes m})}. \quad (76)$$

Taking the logarithm, we get

$$k \geq \log_{\Lambda^+(\rho)} [p(1 - \varepsilon)F(\psi)^{-m}], \quad (77)$$

which is precisely Eq. (73). Alternatively, if we use $\log \Lambda^+$ instead of Λ^+ in the above derivation (noting that $\log \Lambda^+$ also decreases on average under stabilizer protocols due to concavity of the logarithm), we obtain the bound in Eq. (74). ■

Another bound of this kind, which also explicitly depends on $\Lambda^+(\rho)$ and $F(\psi)$ but exhibits a different scaling with respect to ε , was recently obtained in Ref. [47]. We compare the performance of the bounds in Fig. 4.

When $p = 1$, Eq. (76) recovers a related recent bound of Ref. [66]. When $\varepsilon = 0$, we obtain a benchmark on the performance of all distillation protocols that distill the target exactly, but can fail with a certain probability:

$$\frac{k}{mp} \geq \frac{\log F(\psi)^{-1}}{\log \Lambda^+(\rho)}. \quad (78)$$

This was considered for odd-dimensional qudits in Refs. [35,64] as the “distillation efficiency.”

Additionally, the ultimate constraints on the convertibility between two states are often characterized in the asymptotic limit, where we are interested in the best achievable rate $R(\rho \rightarrow \psi)$ at which k copies of ρ can be approximately converted to $kR(\rho \rightarrow \psi)$ copies of ψ , with the error ε of this conversion vanishing in the limit

$k \rightarrow \infty$. Using Eq. (74) with $p = 1$, any such rate must satisfy

$$R(\rho \rightarrow \psi) \leq \frac{\log \Lambda^+(\rho)}{\log F(\psi)^{-1}}, \quad (79)$$

which gives a semidefinite programming upper bound on the asymptotic rate of transformation between any state ρ and a pure state ψ of at most three qubits.

States of interest in magic state distillation include $|H\rangle$ and $|F\rangle$ [17]. These states obey a Clifford symmetry in the following sense; we say a state $|\psi\rangle$ is Clifford symmetric if there exists an Abelian subgroup \mathcal{C}_ψ of the Clifford group such that (i) $C|\psi\rangle = |\psi\rangle$ for all $C \in \mathcal{C}_\psi$; and (ii) $|\psi\rangle$ is the unique state with this property up to a global phase. Crucially, any such state has extent equal to the inverse of its stabilizer fidelity [25], so $\xi(\psi) = F(\psi)^{-1}$. When we already know the value of the extent $\xi(\psi)$, we need only to evaluate $\Lambda^+(\rho)$ to determine the bounds in Theorem 13 and in Eq. (79). For instance, for the rate of transformation from any state to a Clifford symmetric state of up to three qubits, we get $R(\rho \rightarrow \psi) \leq \frac{\log \Lambda^+(\rho)}{\log \Lambda^+(\psi)}$. Asymptotic distillation rates of the magic states $|H\rangle$ and $|F\rangle$ are bounded by

$$R(\rho \rightarrow |H\rangle\langle H|) \leq \frac{\log \Lambda^+(\rho)}{\log(4 - 2\sqrt{2})}, \quad (80)$$

$$R(\rho \rightarrow |F\rangle\langle F|) \leq \frac{\log \Lambda^+(\rho)}{\log(3 - \sqrt{3})}, \quad (81)$$

where we use the known values of $\xi(|H\rangle)$ and $\xi(|F\rangle)$ [25,67].

The above can be compared with the recent bounds obtained in Ref. [64] for qudit magic state theory, as our approach similarly yields computable upper bounds on the rates of distillation, although applicable to the fundamentally important case of qubit systems.

We can alternatively show these asymptotic results by using the regularized relative entropy of magic [35] to bound the achievable rates of transformations between states using any stabilizer protocol. Specifically, define $r_\infty(\rho) = \lim_{n \rightarrow \infty} (1/n)r(\rho^{\otimes n})$ where $r(\rho) = \min_{\sigma \in \bar{\mathcal{S}}_n} D(\rho \parallel \sigma)$ and $D(\rho \parallel \sigma) = \text{Tr}(\rho \log \rho) - \text{Tr}(\rho \log \sigma)$ is the quantum relative entropy. Then the ratio $r_\infty(\rho)/r_\infty(\sigma)$ provides a general upper bound on the rate $R(\rho \rightarrow \sigma)$ of the transformation from ρ to σ using stabilizer protocols [35]. This upper bound is achievable whenever the states can be reversibly interconverted [35] or when the set of stabilizer protocols is relaxed to the class of operations that asymptotically preserve the set of stabilizer states [68]. Using the bounds $r(\rho) \leq \log \Lambda^+(\rho)$ for arbitrary states [69] and $r(\psi) \geq -\log F(\psi)$ for pure states [69], we similarly obtain Eq. (79). Notice also that $r(\psi) = \log \Lambda^+(\psi)$ for any Clifford symmetric state, and $r_\infty(\psi) = \log \Lambda^+(\psi)$ for a Clifford symmetric state of at most three qubits.

Finally, we remark that the best known magic state distillation protocols perform many orders of magnitude worse than our best bounds. It remains a considerable challenge to close this gap.

VII. CLASSICAL SIMULATION ALGORITHMS

In this section we introduce three simulation techniques, each associated to one of the magic monotones defined earlier. In Sec. VII A, we generalize quasiprobability-based methods [15,20,21] for estimating Born-rule probabilities or expectation values of bounded observables up to additive error. We use a novel choice of frame consisting of the set of stabilizer dyads and extend quasiprobabilistic techniques to accommodate this choice. By doing so, we are able to reduce the sampling overhead compared to previous qubit quasiprobability simulators [15,21,22], resulting in a runtime proportional to the dyadic negativity squared, $\Lambda(\rho)^2$. In Sec. VII B we describe a simulator, which extends stabilizer-rank methods [24,25], previously only defined for pure states, to arbitrary mixed-state inputs. The algorithm simulates the sampling of bit strings from a quantum circuit (i.e., by measurement of a subset of qubits in the computational basis). We show that the classical distribution we sample from is δ -close in ℓ_1 norm to the quantum distribution, and that under modest assumptions each string is sampled in average time $\mathcal{O}[\Xi(\rho)/\delta^{-3}]$, where Ξ is the mixed-state extent. When an equimagical decomposition is known (recall Definition 3), this becomes the worst-case runtime. This reduces the

runtime by a factor of δ^{-1} compared to the results of Ref. [25].

Finally, in Sec. VII C, we introduce the constrained path simulation technique, which efficiently estimates Pauli expectation values or Born-rule probabilities up to additive error on $\mathcal{E}(\rho)$ for stabilizer channel \mathcal{E} and nonstabilizer state ρ . The technique approximates the magic state ρ with the stabilizer part of a feasible solution to the generalized robustness problem, Eq. (12). Whereas the dyadic frame simulator outputs estimates to arbitrarily high precision but with runtime that grows with $\Lambda(\rho)$, here the estimate is efficiently computed, but with unavoidable additive error lower bounded as $\mathcal{O}[\Lambda^+(\rho)]$, where Λ^+ is the generalized robustness.

A. Dyadic frame simulator

1. Quasiprobability simulators

Before describing our first algorithm, we briefly review the principles of classical simulation using quasiprobabilities. A very general notion of quasiprobability simulation was introduced by Pashayan *et al.* [20]. A specific instance of this type of simulator is defined by fixing a frame, a finite set of operators that forms a basis for the space of Hermitian operators acting on a Hilbert space. This basis need not be orthonormal and can in general be over-complete. For concreteness we consider the algorithm introduced by Howard and Campbell [21], where the frame is the set of pure stabilizer state projectors. We can define the n -qubit stabilizer frame as

$$\mathcal{G}_n = \{|\phi\rangle\langle\phi| : |\phi\rangle \in S_n\}, \quad (82)$$

so that the convex hull of \mathcal{G}_n is precisely $\bar{\mathcal{S}}_n$, the set of mixed stabilizer states. Indeed, \mathcal{G}_n forms an overcomplete basis for the Hermitian operators on \mathbb{C}^{2^n} . It follows that any n -qubit density ρ matrix has at least one decomposition of the form:

$$\rho = \sum_j q_j |\phi_j\rangle\langle\phi_j|, \quad |\phi_j\rangle\langle\phi_j| \in \mathcal{G}_n, \quad \sum_j q_j = 1, \quad (83)$$

with q_j real. Consider the simulation task of estimating the Born-rule probability $\mu = \text{Tr}[\Pi\mathcal{E}(\rho)]$, where Π is a stabilizer projector and \mathcal{E} is an efficiently simulable channel, but ρ is a general mixed magic state. Given a known quasiprobability decomposition as per Eq. (83), we can rewrite

$$\mu = \sum_j q_j \text{Tr}[\Pi\mathcal{E}(|\phi_j\rangle\langle\phi_j|)] = \sum_j \frac{|q_j|}{\|q\|_1} E_j, \quad (84)$$

where $E_j = \|q\|_1 \text{sign}(q_j) \text{Tr}[\Pi\mathcal{E}(|\phi_j\rangle\langle\phi_j|)]$. Now $|q_j|/\|q\|_1$ are non-negative and sum to unity, so form a proper probability distribution. The Howard and Campbell [21]

algorithm goes as follows. First fix a total number of samples M . Then, perform the below.

1. For each integer k from 1 to M , sample index j_k from the distribution $\{|q_j|/\|q\|_1\}$.
2. Compute each $\widehat{E}_k = E_{j_k}$.
3. Output $\widehat{\mu} = (1/M) \sum_k \widehat{E}_k$.

It is clear that since $\text{Tr}[\Pi \mathcal{E}(|\phi_j\rangle\langle\phi_j|)]$ amounts to evaluating a stabilizer circuit, each \widehat{E}_k can be efficiently computed using the standard Gottesman-Knill tableaux method [2,3]. Moreover, one can easily check that $\mathbb{E}(\widehat{\mu}) = \sum_j (|q_j|/\|q\|_1) E_j = \mu$, so the algorithm gives an unbiased estimator for the Born-rule probability. However, due to the renormalization of the distribution, each estimate \widehat{E}_k takes a value in the range $[-\|q\|_1, +\|q\|_1]$, increasing the variance of the estimator. From Hoeffding's inequalities [70], the probability that $\widehat{\mu}$ is far from the expected value μ is bounded as

$$\Pr\{|\widehat{\mu} - \mu| \geq \epsilon\} \leq 2 \exp\left(-\frac{M\epsilon^2}{2\|q\|_1^2}\right). \quad (85)$$

It follows that to estimate the value within additive error at most ϵ with probability at least $1 - p_{\text{fail}}$, we must set the number of samples so that $M \geq 2\|q\|_1^2 \epsilon^{-2} \log(2p_{\text{fail}}^{-1})$. Recall from Definition 4 that robustness of magic $\mathcal{R}(\rho)$ is defined as the minimal $\|q\|_1$, so the worst-case runtime for the Howard and Campbell algorithm scales with (at least) $\mathcal{R}(\rho)^2$.

Whereas in the simulation model described above, the frame is comprised of stabilizer projectors $|\phi\rangle\langle\phi|$, in our dyadic frame simulator we extend the frame to include dyads $|L\rangle\langle R|$ where $|L\rangle$ and $|R\rangle$ may be different stabilizer states. An operator is now considered free if it is in the convex hull of the dyads $e^{i\theta}|L\rangle\langle R|$. Importantly, a density matrix σ can be written in this form if and only if $\sigma \in \widehat{S}_n$. Nonfree density matrices are then expressed as generalized quasiprobability distributions over the set of n -qubit dyads, where the ‘‘quasiprobabilities’’ are now complex valued. As we shall see, the associated dyadic negativity quantifies the classical simulation overhead for estimating Born-rule probabilities on a nonfree state. In the next subsection we illustrate our new algorithm by giving a simplified version where the stabilizer circuit elements are restricted to be probabilistic mixtures of Clifford gates. We subsequently generalize the algorithm to cover all completely stabilizer-preserving circuits with magic state inputs.

2. Dyadic frame simulator

We assume the following restricted simulation setting. The input to the algorithm will consist of (i) a known dyadic decomposition of a mixed magic state $\rho = \sum_j \alpha_j |L_j\rangle\langle R_j|$; (ii) a circuit description comprised of a list of T quantum operations $\{O^{(1)}, \dots, O^{(T)}\}$; and (iii)

a stabilizer projector Π representing the outcome of a Pauli measurement. We stipulate that each $O^{(t)}$ must be a convex mixture of unitary Clifford channels, $O^{(t)} = \sum_k p_k^{(t)} U_k(\cdot) U_k^\dagger$, and we assume this decomposition is known and can be efficiently sampled from. The output of the algorithm is again an estimate for the Born-rule probability $\mu = \text{Tr}[\Pi \mathcal{E}(\rho)]$, where $\mathcal{E} = O^{(T)} \circ \dots \circ O^{(1)}$. Note that the above restriction on $O^{(T)}$ means that we can write the whole circuit as an ensemble over unitary Clifford gates:

$$\mathcal{E}(\cdot) = \sum_{\vec{k}} p_{\vec{k}} U_{\vec{k}}(\cdot) U_{\vec{k}}^\dagger, \quad (86)$$

where $\vec{k} = (k_1, k_2, \dots, k_T)$ is a vector that represents a Clifford trajectory through the circuit $U_{\vec{k}} = U_{k_T} \dots U_{k_2} U_{k_1}$, and $p_{\vec{k}}$ is a product distribution and so can be efficiently sampled from. The algorithm proceeds by sampling elements from the initial distribution, computing an estimate, repeating many times, and averaging. The procedure for generating one sample is as follows.

1. Randomly select index j with probability $|\alpha_j|/\|\alpha\|_1$.
2. Randomly select trajectory \vec{k} with probability $p_{\vec{k}}$.
3. Compute final dyad:

$$e^{i\theta'_{j,\vec{k}}} |L'_{j,\vec{k}}\rangle\langle R'_{j,\vec{k}}| = e^{i\theta'_j} U_{\vec{k}} |L_j\rangle\langle R_j| U_{\vec{k}}^\dagger. \quad (87)$$

4. Compute sample $\widehat{E} = \|\alpha\|_1 \text{Re}\{e^{i\theta'_{j,\vec{k}}} \langle R'_{j,\vec{k}} | \Pi | L'_{j,\vec{k}} \rangle\}$.

In step 3, $e^{i\theta'}$ is a final global phase taking into account the initial phase $e^{i\theta_j} = \alpha_j/|\alpha_j|$ and the action of the sampled unitary circuit on $|L_j\rangle$ and $|R_j\rangle$, respectively. Whereas the Howard and Campbell algorithm dealt with projectors $|\phi\rangle\langle\phi|$, so that any global phase on $|\phi\rangle$ is unimportant, here $|L_j\rangle$ and $|R_j\rangle$ can represent different stabilizer states and the combined phase can affect both the magnitude and sign of the real-valued sample \widehat{E} . While the original tableaux method used in the Gottesman-Knill theorem does not track this global phase, subsequent extensions of the method show that the update can be efficiently computed, including the phase [23–25]. We can also efficiently compute the complex inner product $\langle L|R\rangle$ for any pair of stabilizer states [23–25]. Thus steps 3 and 4 are efficient. Note that the two parts of the dyad $U_{\vec{k}} |L_j\rangle$ and $\langle R_j| U_{\vec{k}}^\dagger = (U_{\vec{k}} |R_j\rangle)^\dagger$ are updated independently.

The algorithm is completed by repeating steps 1–4 M times. We can check that the method gives an unbiased estimator for the target Born-rule probability:

$$\mathbb{E}(\widehat{E}) = \sum_{j,\vec{k}} \frac{|\alpha_j|}{\|\alpha\|_1} p_{\vec{k}} \left(\|\alpha\|_1 \text{Re}\{e^{i\theta'_{j,\vec{k}}} \langle R'_{j,\vec{k}} | \Pi | L'_{j,\vec{k}} \rangle\} \right) \quad (88)$$

$$= \text{Re} \left\{ \sum_{j, \vec{k}} e^{i\theta_j} |\alpha_j| p_{\vec{k}} \text{Tr}[\Pi U_{\vec{k}} |L_j\rangle\langle R_j| U_{\vec{k}}^\dagger] \right\} \quad (89)$$

$$= \text{Re} \left\{ \text{Tr} \left[\Pi \sum_{\vec{k}} p_{\vec{k}} U_{\vec{k}} \left(\sum_j \alpha_j |L_j\rangle\langle R_j| \right) U_{\vec{k}}^\dagger \right] \right\} \quad (90)$$

$$= \text{Tr}[\Pi \mathcal{E}(\rho)]. \quad (91)$$

We can therefore apply Hoeffding's inequality in the same way as for the standard quasiprobability technique, and using the fact that each \tilde{E} is in the range $[-\|\alpha\|_1, +\|\alpha\|_1]$, we find that the total number of samples needed to achieve additive error ϵ and success probability $1 - p_{\text{fail}}$ is

$$M \geq 2\|\alpha\|_1^2 \epsilon^{-2} \log(2p_{\text{fail}}^{-1}). \quad (92)$$

When the decomposition of ρ is optimal with respect to dyadic negativity as per Definition 5, we have that $\|\alpha\|_1 = \Lambda(\rho)$. When this holds, the worst-case runtime of the algorithm will be $\mathcal{O}[\Lambda(\rho)^2]$.

This simplified algorithm can be used only in the case where the stabilizer circuit is a convex mixture of unitary Clifford operations, so channels are restricted to be unital. Our main goal, however, is to admit more general stabilizer channels. In particular, extending to adaptive Clifford circuits with mixed magic state inputs allows for universal quantum computation [17]. We now sketch how the dyadic frame simulator can be extended to admit all completely stabilizer-preserving channels. Full pseudocode and technical proofs of validity and performance are given in Appendix C.

The simplicity of the restricted simulator derives from the fact that unitary operations preserve the norm of the state vector. This means that when each circuit element $O^{(i)}$ can be decomposed as a convex mixture of unitary gates, the probability of choosing a particular trajectory \vec{k} through the circuit depends only on the coefficients $p_{\vec{k}}^{(i)}$ and is independent of initial state. Conversely, Kraus decompositions of nonunitary channels always include nonunitary operators. When these channels appear in a circuit, transition probabilities for selecting one of the nonunitary operators must be computed on the fly as we step through the circuit. These transition probabilities depend not only on the initial state, but on the Kraus operators selected in previous steps, so they cannot be precomputed. Note that in general a channel may be decomposed as a mixture of N_U unitary and N_K nonunitary Kraus operators:

$$O(\cdot) = \sum_k^{N_U} p_k U_k(\cdot) U_k^\dagger + \sum_{k'}^{N_K} q_{k'} K_{k'}(\cdot) K_{k'}^\dagger. \quad (93)$$

The probability of picking one of the N_U operators U_k can simply be read off from the coefficients p_k . We can infer the

total probability $1 - \sum_k p_k$ that the trajectory chosen will be from among the N_K nonunitary operators, but individual transition probabilities for each $K_{k'}$ must be computed based on the initial state. In Appendix C we show how appropriate transition probabilities can be computed efficiently even when the input operator is not a state but a dyad, provided that we restrict to channel decompositions where $N_U, N_K \leq \text{poly}(n)$. We call such a decomposition *simulable*. This leads to the following theorem.

Theorem 14. *Let $\rho = \sum_j \alpha_j |L_j\rangle\langle R_j|$, be a known dyadic decomposition of an initial n -qubit state, where $\alpha \in \mathbb{C}$ and the probability distribution $\{|\alpha_j|/\|\alpha\|_1\}$ can be efficiently sampled. Let $\mathcal{E} = O^{(T)} \circ \dots \circ O^{(1)}$, where each $O^{(i)} \in \mathcal{O}_n$ is a completely stabilizer-preserving channel. Suppose that every $O^{(i)}$ has a known simulable decomposition. Then, given a stabilizer projector Π , we can estimate the Born-rule probability $\mu = \text{Tr}(\Pi \mathcal{E}[\rho])$ within additive error ϵ , with success probability at least $1 - p_{\text{fail}}$ and worst-case runtime*

$$\frac{\|\alpha\|_1^2}{\epsilon^2} \log(p_{\text{fail}}^{-1}) T \text{poly}(n). \quad (94)$$

Furthermore, if the dyadic decomposition of ρ is optimal then $\|\alpha\|_1$ can be replaced by $\Lambda(\rho)$.

By exploiting a dyadic frame, the negativity of the quasiprobability distribution and algorithm runtime is greatly reduced compared to previous work [21], with an improved exponential scaling of the runtime (recall Theorem 12).

B. The density-operator stabilizer-rank simulator

1. Prior art: the BBCCGH simulator

Here we briefly review a previous stabilizer rank-based simulation method, which we refer to as BBCCGH in what follows (after the authors' initials [25]). BBCCGH simulates sampling length w bit strings \vec{x} from measurements on pure magic states $|\psi\rangle$ with runtime linear in pure-state extent $\xi(|\psi\rangle)$, and represents the prior state of the art in stabilizer-rank techniques. In subsequent sections we improve on this algorithm and generalize to mixed states. BBCCGH can be decomposed into two main subroutines: SPARSIFY, which generates a sparse approximation of the target state, and FASTNORM, which estimates Born-rule probabilities $\|\Pi|\psi\rangle\|^2$ up to multiplicative error. By calling FASTNORM $\mathcal{O}(w)$ times, one estimates a chain of conditional probabilities so as to successively sample the outcome for each bit of \vec{x} in turn. It is crucial that the error is multiplicative, as this ensures that the output distribution of the classical algorithm is close in ℓ_1 norm to the quantum distribution $P(\vec{x}) = |\langle \vec{x} | \psi \rangle|^2$.

In general, stabilizer-rank simulators exploit the fact that any pure quantum state $|\psi\rangle$ can be expressed as a linear

combination of stabilizer states,

$$|\psi\rangle = \sum_{j=1}^k c_j |\phi_j\rangle, \quad (95)$$

where $|\phi_j\rangle$ are stabilizer states and c_j are complex. The exact stabilizer rank $\chi(|\psi\rangle)$ is the smallest number of terms k needed for a given state $|\psi\rangle$ [23–25]. Computations can be performed in $\text{poly}(n, k)$ time by treating each stabilizer term in turn (albeit k can grow exponentially with n). In particular Bravyi *et al.* [25] showed that FASTNORM can estimate $\|\psi\|^2$ up to multiplicative error by repeatedly generating a random number $\eta_A = 2^n |\langle \phi_A | \psi \rangle|^2$, where $|\phi_A\rangle$ is randomly drawn from a subset of stabilizer states known as equatorial states. Evaluating η_A amounts to computing k stabilizer inner products (one for each term of ψ), which can be done efficiently by exploiting a canonical representation of stabilizer states known as CH form [25]. We summarize this result of Bravyi *et al.* in the following theorem.

Theorem 15 ([25]). *Given an unnormalized n -qubit vector $|\psi\rangle = \sum_{j=1}^k c_j |\phi_j\rangle$ with χ stabilizer terms in its decomposition, there exists a classical algorithm FASTNORM that outputs a random variable η such that*

$$(1 - \epsilon)\|\psi\|^2 \leq \eta \leq (1 + \epsilon)\|\psi\|^2 \quad (96)$$

with probability greater than $1 - p_{\text{fail}}$ in worst-case runtime $\mathcal{O}[\chi n^3 \epsilon^{-2} \log(p_{\text{fail}}^{-1})]$.

By applying this algorithm to projected vectors $\Pi|\psi\rangle$, where Π is a stabilizer projector, one can estimate Born-rule probabilities $\|\Pi|\psi\rangle\|^2$.

If one was to apply FASTNORM directly to the ideal state $|\psi\rangle$, the runtime would be $\mathcal{O}[\chi(\psi)]$, where χ is the exact stabilizer rank. However, computing the exact stabilizer rank is intractable for many-qubit states. Instead the strategy of BBCCGH is to approximate $|\psi\rangle$ with a sparsified k -term vector $|\Omega\rangle$ of smaller stabilizer rank, using the subroutine SPARSIFY (Fig. 5). Bravyi *et al.* [25, Lemma 6] showed that for any pure state $|\psi\rangle$ and any integer $k > 0$, one can use SPARSIFY to generate random (unnormalized) states $|\Omega\rangle$ with k stabilizer terms such that

$$\mathbb{E}(\| |\psi\rangle - |\Omega\rangle \|^2) \leq \frac{\|c\|_1^2}{k}. \quad (97)$$

In Appendix D we present a simple corollary of Ref. [25, Lemma 6], which implies that

$$\mathbb{E}(\| |\psi\rangle\langle\psi| - |\Omega\rangle\langle\Omega| \|_1) \leq 2 \frac{\|c\|_1}{\sqrt{k}} + \frac{\|c\|_1^2}{k} \approx 2 \frac{\|c\|_1}{\sqrt{k}}. \quad (98)$$

For any target precision $\delta_S > 0$, choosing k so that

$$k \geq \frac{4\|c\|_1^2}{\delta_S^2}, \quad (99)$$

we get

$$\mathbb{E}(\| |\psi\rangle\langle\psi| - |\Omega\rangle\langle\Omega| \|_1) \leq \delta_S + \mathcal{O}(\delta_S^2). \quad (100)$$

Recall from Sec. II that the minimal value of $\|c\|_1^2$ is precisely the pure-state extent. We call Eq. (100) combined with the lower bound on k in Eq. (99) the BBCCGH sparsification lemma [25]. Thus, with high probability and subject to some technical caveats discussed in Appendix E, by combining the two subroutines BBCCGH simulates sampling from the quantum distribution $P(\vec{x}) = |\langle \vec{x} | \psi \rangle|^2$ up to an error δ_S in runtime $\|c\|_1^2 \delta_S^{-4} \text{poly}(n, w)$. Assuming an optimal decomposition [$\xi(\psi) = \|c\|_1^2$], the runtime therefore scales linearly with pure-state extent.

Below we improve on this algorithm in three main respects: (i) we extend the simulator from pure to mixed magic state inputs, so that the average-case runtime is proportional to the mixed-state extent Ξ defined in Sec. II; (ii) we show that important cases admit decompositions such that Ξ quantifies the worst-case runtime; and (iii) we derive a new sparsification lemma that improves the runtime over that implied by Eq. (100) by a factor of $1/\delta_S$ with minor caveats. Our new sparsification lemma also avoids some technical difficulties that arise when applying the

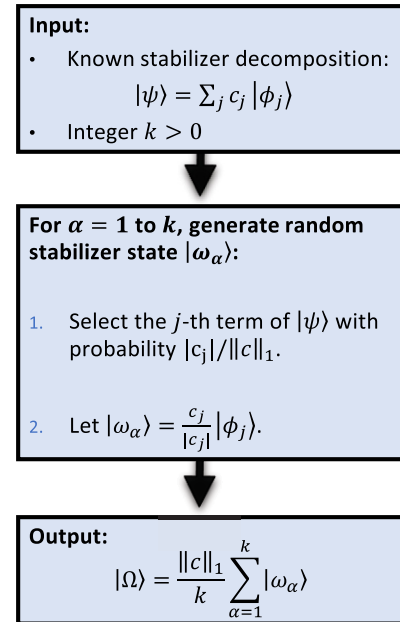


FIG. 5. The SPARSIFY procedure introduced by Bravyi *et al.* [25]. The exact state $|\psi\rangle$ is approximated by an unnormalized random vector $|\Omega\rangle$ with k stabilizer terms that is, on average, δ_S^2 -close in the Euclidean norm, where $\delta_S = \|c\|_1^2/k$.

BBCCGH sparsification lemma in a practical algorithm. The runtime improvements originate from working in the density-operator picture even when the input magic state is pure.

We first discuss the proof of our new lemma, before applying it to classically simulate bit-string sampling. While our ideas naturally apply to estimating Born probabilities, and can be extended to propagate an initial state through a noisy stabilizer circuit prior to measurement, we omit this for brevity.

2. Sparsification lemma

The input to the subroutine SPARSIFY is an integer k and pure state $|\psi\rangle$ with known stabilizer decomposition, Eq. (95), with coefficient vector c . The output is a randomly chosen k -term sparsification of $|\psi\rangle$,

$$|\Omega\rangle = \frac{\|c\|_1}{k} \sum_{\alpha=1}^k |\omega_\alpha\rangle, \quad (101)$$

where each $|\omega_\alpha\rangle$ is an independent and identically distributed (i.i.d.) sampled stabilizer state $(c_j/|c_j|)|\phi_j\rangle$ for some j (see Fig. 5), so that we have [25]

$$\mathbb{E}(|\omega_\alpha\rangle) = \frac{|\psi\rangle}{\|c\|_1} \Rightarrow \mathbb{E}(|\Omega\rangle) = |\psi\rangle. \quad (102)$$

Since the output $|\Omega\rangle$ of SPARSIFY is a random superposition of nonorthogonal terms, it need not have unit norm. In Ref. [25], after obtaining a state $|\Omega\rangle$ from SPARSIFY, one estimates its Euclidean norm, and discards the state if its norm is not close to 1. A state postselected in this way will be close to the target state with high probability. See Appendix E for a discussion of why this postselection is necessary.

Here, we instead consider a sampling strategy that avoids postselecting $|\Omega\rangle$. After SPARSIFY gives a random $|\Omega\rangle$, we renormalize so that it has unit norm. Furthermore, instead of bounding the error between an individual sample and the target state $|\psi\rangle$, we bound the error between $|\psi\rangle$ and the whole ensemble as captured by the density matrix

$$\rho_1 := \mathbb{E} \left[\frac{|\Omega\rangle\langle\Omega|}{\langle\Omega|\Omega\rangle} \right] = \sum_{\Omega} \Pr(\Omega) \frac{|\Omega\rangle\langle\Omega|}{\langle\Omega|\Omega\rangle}. \quad (103)$$

Intuitively, this is advantageous because coherent errors in each sample smooth out to a less harmful stochastic error. Similarly, randomizing coherent errors improves error bounds in the setting of circuit compilation [71–74].

Our refinement to the BBCCGH sparsification lemma is summarized in the following theorem.

Theorem 16. *Let ρ_1 be the mixed state in Eq. (103). Let $|\psi\rangle$ be an input state with known decomposition $|\psi\rangle =$*

$\sum_j c_j |\phi_j\rangle$, where $|\phi_j\rangle$ are stabilizer states, and let c be the vector whose elements are the coefficients c_j and

$$C_{\psi,c} = \|c\|_1 \sum_j |c_j| |\langle\psi|\phi_j\rangle|^2. \quad (104)$$

Then there is a critical precision $\delta_c = 8(C_{\psi,c} - 1)/\|c\|_1^2$ such that for every target precision δ_S for which $\delta_S \geq \delta_c$, we can sample pure states from an ensemble ρ_1 , where every pure state drawn from ρ_1 has stabilizer rank at most $\lceil 4\|c\|_1^2/\delta_S \rceil$ and

$$\|\rho_1 - |\psi\rangle\langle\psi|\|_1 \leq \delta_S + \mathcal{O}(\delta_S^2). \quad (105)$$

When $|\psi\rangle$ is a Clifford magic state (recall Sec. VI), the critical precision becomes $\delta_c = 0$, and sampled pure states in ρ_1 have stabilizer rank at most $\lceil (2 + \sqrt{2})\|c\|_1^2/\delta_S \rceil$.

Notice that the theorem sets a critical precision δ_c above which we can achieve the promised $1/\delta_S$ improvement in the runtime over BBCCGH [25]. At higher precision, our runtime has the same leading-order δ_S scaling as BBCCGH but with a much smaller constant prefactor, so still yields improved performance. Furthermore, for the important case of noisy T states, they are Clifford magic states so the improvement holds across all δ .

The proof of Theorem 16 follows from two lemmata. Here we sketch the proof strategy, deferring full technical proofs to Appendices F and G. The first lemma captures the idea that the ensemble, Eq. (103), can be made close in the trace norm to the target state $|\psi\rangle\langle\psi|$ by choosing sufficiently large k , up to a term that depends on the variance of $\langle\Omega|\Omega\rangle$. The second lemma then bounds this variance in terms of $C_{\psi,c}$, $\|c\|_1$, and k .

Lemma 17 (Ensemble sampling lemma). *Given a state $|\psi\rangle = \sum_j c_j |\phi_j\rangle$ where ϕ_j are stabilizer states, we can sample from an ensemble ρ_1 such that every sampled pure state has stabilizer rank $\leq k$ and*

$$\|\rho_1 - |\psi\rangle\langle\psi|\|_1 \leq \frac{2\|c\|_1^2}{k} + \sqrt{\text{Var}[\langle\Omega|\Omega\rangle]}, \quad (106)$$

where $|\Omega\rangle$ is the random sparsified vector defined in Eq. (101).

The first step in proving Lemma 17 is to note that we can use the triangle inequality to split the problem into two parts:

$$\|\rho_1 - |\psi\rangle\langle\psi|\|_1 \leq \|\rho_1 - \rho_2\|_1 + \|\rho_2 - |\psi\rangle\langle\psi|\|_1, \quad (107)$$

where $\rho_2 = \mathbb{E}(|\Omega\rangle\langle\Omega|)/\mathbb{E}(\langle\Omega|\Omega\rangle)$. The first term is upper bounded by $\text{Var}[\langle\Omega|\Omega\rangle]$. The second term can then be evaluated in terms of $|\omega\rangle$, and turns out to be upper bounded by

$2\|c\|_1^2/k$. Full technical details are given in Appendix F. It remains to bound the variance of $\langle\Omega|\Omega\rangle$.

Lemma 18 (Sparsification variance bound). *Using the notation of Lemma 17 the variance of $\langle\Omega|\Omega\rangle$ satisfies the bound*

$$\text{Var}[\langle\Omega|\Omega\rangle] \leq \frac{4(C-1)}{k} + \frac{2\|c\|_1^4}{k^2} + \mathcal{O}\left(\frac{C}{k^3}\right), \quad (108)$$

where $C = C_{\psi,c}$ is as given in Eq. (104). When $|\psi\rangle$ is a Clifford magic state as defined in Ref. [25],

$$\text{Var}[\langle\Omega|\Omega\rangle] \leq \frac{2\|c\|_1^4}{k^2} + \mathcal{O}\left(\frac{1}{k^3}\right). \quad (109)$$

In Appendix G, we prove Lemma 18 by expanding $\text{Var}[\langle\Omega|\Omega\rangle]$ as a series of terms of the form $\mathbb{E}(\langle\omega_\alpha|\omega_\beta\rangle\langle\omega_\lambda|\omega_\mu\rangle)$, treating the cases where the indices $\alpha, \beta, \lambda,$ and μ are all distinct (and therefore correspond to i.i.d. random variables), where $\alpha = \beta$, but (α, λ, μ) are all distinct and so on.

By combining Lemmas 17 and 18 we can now prove Theorem 16. Substituting $k = 4\|c\|_1^2/\delta_S$ and $\delta_S \geq 8(C-1)/\|c\|_1^2$ into Eq. (108), we obtain

$$\text{Var}[\langle\Omega|\Omega\rangle] \leq \frac{\delta_S^2}{4} \left[1 + \mathcal{O}\left(\frac{\delta_S}{\|c\|_1^4}\right) \right], \quad (110)$$

and hence, using $\sqrt{1+x} \leq 1+x$ for $x \geq 0$:

$$\sqrt{\text{Var}[\langle\Omega|\Omega\rangle]} \leq \frac{\delta_S}{2} + \mathcal{O}(\delta_S^2). \quad (111)$$

Using Eq. (111) with the expression for k and Lemma 17, we have

$$\|\rho_1 - |\psi\rangle\langle\psi|\|_1 \leq \delta_S + \mathcal{O}(\delta_S^2). \quad (112)$$

This proves the main result of Theorem 16. When $|\psi\rangle$ is a Clifford magic state, Eq. (109) combined with Lemma 17 gives

$$\|\rho_1 - |\psi\rangle\langle\psi|\|_1 \leq \frac{(2 + \sqrt{2})\|c\|_1^2}{k} + \mathcal{O}\left(\frac{1}{k^2}\right). \quad (113)$$

This allows us to obtain Eq. (105) by setting $k = \lceil(2 + \sqrt{2})\|c\|_1^2/\delta_S\rceil$, completing the proof.

We showed that whenever the constraint on the target precision δ_S is greater than a critical precision, one can sample from an ensemble of sparsified states ρ_1 that is δ_S -close in the trace norm to $\langle\psi|\psi\rangle$, where the number of stabilizer terms is $k = 4\|c\|_1^2/\delta_S$. Compared to the BBCCGH [25] sparsification lemma where $k = 4\|c\|_1^2/\delta_S^2$, we see a factor $1/\delta_S$ improvement. If the target precision is smaller

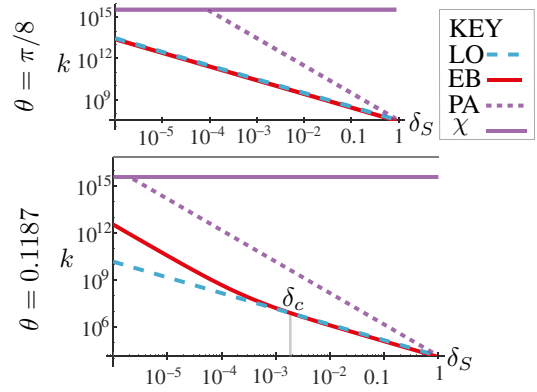


FIG. 6. For the target state $|\psi\rangle = [\cos(\theta)|0\rangle + \sin(\theta)|1\rangle]^{\otimes 100}$ with two choices of θ , we plot the trace norm error δ_S when using a sparsification with k terms. Exact bound (EB) refers to Eq. (106) and is valid for all δ_S , with the variance exactly bounded by Eq. (G33). Leading order (LO) refers to our Theorem 16 expression $k = 4\|c\|_1^2/\delta_S$, and is valid provided $\delta_S \geq \delta_c$ with δ_c highlighted by a vertical line. Note $\theta = \pi/8$ corresponds to the Clifford magic state $|H\rangle$, for which $\delta_c = 0$. Prior art (PA) shows the cost of Ref. [25]. The exact stabilizer rank is χ (see Theorem 2 of Ref. [25]) and this is an upper bound on PA. When $C \neq 1$ and $\delta_S < \delta_c$, then EB shows that there is still a large saving even though LO is not valid in this regime. To better understand the deviations of EB from LO, we refer the reader to Appendix G and in particular Fig. 8, and to the discussion at the end of Sec. VII B.

than the critical precision, one can compute C and obtain a sharp bound on the trace-norm error by using Lemmas 17 and 18 directly. In this case, the δ_S^{-2} scaling of k is recovered, but with a prefactor often much smaller than in the original BBCCGH sparsification lemma. This is because one typically finds that $(C-1)/\|c\|_1^2 \ll 1$ for many-qubit magic states. We illustrate this in Fig. 6, where we compare the sharpened trace-norm bound of our lemmata with that of Ref. [25] for states of the form $|\psi_N\rangle = |\psi\rangle^{\otimes N}$, where $|\psi\rangle$ are single-qubit magic states, and $N = 100$. While $\delta_S \geq 8(C-1)/\|c\|_1^2$ we have a quadratic improvement over Eq. (98), but even in the high-precision regime, we find a significant reduction in k .

3. Bit-string sampling from mixed magic states

Consider the setting where we have an n -qubit mixed magic state ρ , and we measure a subset of w qubits in the computational basis (i.e., we measure Pauli Z for each qubit), thereby generating a random bit string \vec{x} of length w representing the measurement outcomes. Without loss of generality we can assume we measure the first w qubits. Let $\Pi_{\vec{x}} = |\vec{x}\rangle\langle\vec{x}| \otimes \mathbb{1}_{n-w}$ be the projector representing the outcome where we obtain bit string \vec{x} . Then the probability of obtaining the string \vec{x} is given by the Born rule:

$$P(\vec{x}) = \text{Tr}[\Pi_{\vec{x}}\rho]. \quad (114)$$

We call P the *quantum* probability distribution. Here we deal with the simulation task of classically sampling from a probability distribution $P_{\text{sim}}(\vec{x})$ over w bit strings \vec{x} such that P_{sim} is δ -close in ℓ_1 norm to P , with high probability. Our algorithm is closely related to the sampling algorithm given in Bravyi *et al.* [25], differing in two key respects: (i) whereas the Bravyi *et al.* simulator is defined only for pure states, our variant admits general mixed states; and (ii) we take advantage of our improved sparsification lemma to reduce runtime. We also avoid a postselection step needed for the Bravyi *et al.* algorithm (see Appendix E). Figure 7 gives the key steps for our procedure for sampling a single bit string. Full pseudocode is given in Appendix H. The main steps in the algorithm are (1) the sampling of a random pure state $|\psi_j\rangle$ from the ensemble $\rho = \sum_j p_j |\psi_j\rangle\langle\psi_j|$, (2) a call to the subroutine SPARSIFY to generate the k -term approximation $|\Omega\rangle$, and (3) computation of a chain of conditional probabilities using at most $2w + 1$ calls to FASTNORM. In Appendix H we prove the validity of the algorithm, and give a full analysis of the runtime. Here we first sketch the proof before discussing the runtime improvement over Bravyi *et al.* [25]. In what follows we assume that $\delta > \delta_c$ as defined in Eq. (16), returning to the case of arbitrary precision at the end of the section.

We want to show that the classical probability distribution P_{sim} satisfies

$$\|P_{\text{sim}} - P\|_1 \leq \delta + \mathcal{O}(\delta^2), \quad (115)$$

where P is the quantum distribution. We split the proof into two parts. First, we consider an idealized algorithm EXACT where the calls to FASTNORM are replaced by an oracle that can compute $\|\Pi_{\vec{x}}|\Omega\rangle\|$ exactly given k -term sparsification $|\Omega\rangle$. Let $P_{\text{ex}}(\vec{x})$ be the probability of obtaining the string \vec{x} as the output of EXACT. We first show that P_{ex} is δ_S -close to the quantum distribution P in ℓ_1 norm, and then show that P_{sim} is ϵ -close to P_{ex} . We then split the error budget so that $\delta = \delta_S + \epsilon$. In Appendix H we show that the optimal strategy is to set $\delta_S = \delta/3$ and $\epsilon = 2\delta/3$.

Let $\vec{x}_m = (x_1, \dots, x_m)$ be the bit string comprised of the first m bits of \vec{x} , and let $|\Omega_m\rangle = \Pi_{\vec{x}_m}|\Omega\rangle$ be the projection of the first m qubits of $|\Omega\rangle$. By inspection of the last two steps of Fig. 7, we can multiply the chain of conditional probabilities and obtain the probability of sampling \vec{x} from EXACT given fixed sparsification $|\Omega\rangle$:

$$\Pr(\vec{x}|\Omega) = \Pr(x_1)\Pr(x_2|\vec{x}_1) \dots \Pr(x_w|\vec{x}_{w-1}), \quad (116)$$

$$= \frac{\|\Omega_1\|^2 \|\Omega_2\|^2}{\|\Omega\|^2 \|\Omega_1\|^2} \dots \frac{\|\Pi_{\vec{x}}|\Omega\rangle\|^2}{\|\Omega_{w-1}\|^2}, \quad (117)$$

$$= \frac{\|\Pi_{\vec{x}}|\Omega\rangle\|^2}{\|\Omega\|^2} = \text{Tr} \left[\Pi_{\vec{x}} \frac{|\Omega\rangle\langle\Omega|}{\langle\Omega|\Omega\rangle} \right]. \quad (118)$$

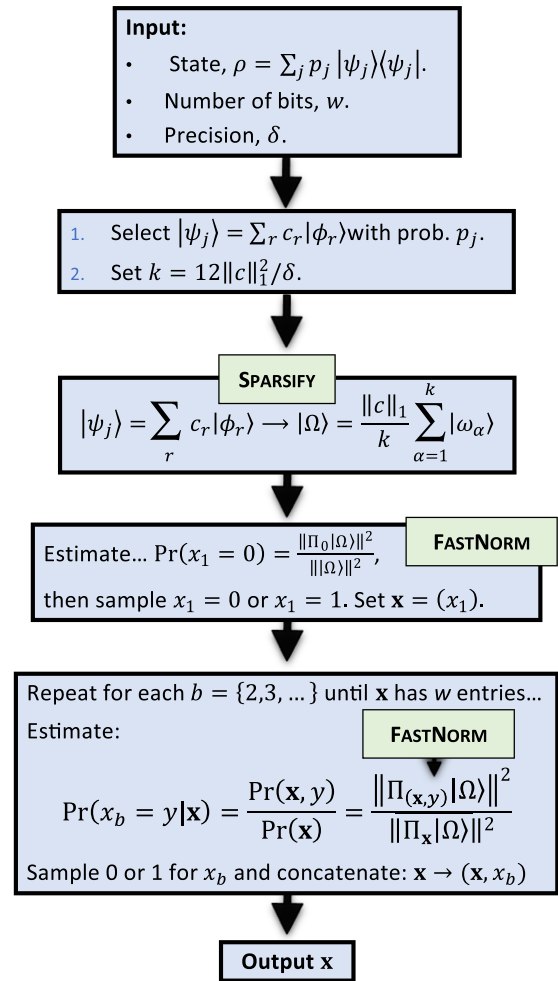


FIG. 7. Our procedure for classically sampling a single length w bit string given an n -qubit state with known decomposition $\rho = \sum_j p_j |\psi_j\rangle\langle\psi_j|$, where each pure state $|\psi_j\rangle$ in turn has a known stabilizer decomposition $|\psi_j\rangle = \sum_r c_r^{(j)} |\phi_r\rangle$. We assume that δ is greater than the critical precision defined in Theorem 16. The procedure is a variant of that given in Ref. [25] for pure-state input, making use of two subroutines from that work, SPARSIFY and FASTNORM, as described above. In the main text we describe how we improve on the sparsification step, and extend the simulator to admit mixed states as input. The factor of 12 in the initial step arises from optimization of the error budget between the sparsification error δ_S and fast norm estimation error ϵ ; we set $\delta_S = \delta/3$ and $\epsilon = 2\delta/3$ (see Appendix H).

Thus EXACT simulates sampling from the quantum state $|\Omega\rangle/\|\Omega\|$ exactly; any error arises solely from the sparsification procedure. Now consider that randomly choosing a pure state ψ_j from $\rho = \sum_j p_j |\psi_j\rangle\langle\psi_j|$, generating a random approximation $|\Omega\rangle$ using SPARSIFY and then normalizing is equivalent to sampling a pure state from the ensemble:

$$\sigma = \sum_j p_j \sum_{\Omega} \Pr(\Omega|\psi_j) \frac{|\Omega\rangle\langle\Omega|}{\langle\Omega|\Omega\rangle} = \sum_j p_j \rho_1^{(j)}, \quad (119)$$

where $\Pr(\Omega|\psi_j)$ is the probability of SPARSIFY outputting the vector $|\Omega\rangle$, and $\rho_1^{(j)}$ is the expected projector $\mathbb{E}(|\Omega\rangle\langle\Omega|/\langle\Omega|\Omega\rangle)$ as defined in Eq. (103), both conditioned on the input to SPARSIFY being $|\psi_j\rangle$. From our argument above it follows that $P_{\text{ex}}(\vec{x}) = \text{Tr}[\Pi_{\vec{x}}\sigma]$. A key conceptual difference between our method and that of Bravyi *et al.* [25] is that while the BBCCGH sparsification results are concerned with the fidelity between the target state $|\psi\rangle$ and a single randomly chosen sparsification $|\Omega\rangle$, here we compare the target state ρ with the full ensemble over sparsifications σ . From our sparsification lemma (Theorem 16), for each pure state $|\psi_j\rangle$, we have that $\|\rho_1^{(j)} - |\psi_j\rangle\langle\psi_j|\|_1 \leq \delta_S + \mathcal{O}(\delta_S^2)$. It follows that $\|\sigma - \rho\|_1 \leq \delta_S + \mathcal{O}(\delta_S^2)$, and so

$$\|P_{\text{ex}} - P\|_1 \leq \delta_S + \mathcal{O}(\delta_S^2). \quad (120)$$

Next we argue that P_{ex} is ϵ -close to P_{sim} , the distribution arising from our full classical algorithm. Recall from Theorem 15 that FASTNORM is able to output estimates for $\|\Omega_m\|^2$ up to some relative error ϵ_{FN} , which we can set arbitrarily small (at the cost of increased runtime). One can show (see Appendix H) that estimating the chain of w conditional probabilities, Eq. (117), using FASTNORM leads to a total relative error $3w\epsilon_{\text{FN}}$ in the distribution sampled from, i.e.,

$$(1 - 3w\epsilon_{\text{FN}})P_{\text{ex}}(\vec{x}) \leq P_{\text{sim}}(\vec{x}) \leq (1 + 3w\epsilon_{\text{FN}})P_{\text{ex}}(\vec{x}), \quad (121)$$

so to achieve relative error ϵ we must set $\epsilon_{\text{FN}} = \epsilon/3w$. This governs the runtime of FASTNORM. By combining this result with Eq. (120) we have

$$\|P_{\text{sim}} - P\|_1 \leq \delta + \mathcal{O}(\delta^2). \quad (122)$$

To analyze the runtime of our simulator, we define

$$\tilde{\Xi} = \sum_j p_j \|c^{(j)}\|_1^2, \quad (123)$$

where $c^{(j)}$ is the vector of coefficients in the decomposition $|\psi_j\rangle = \sum_r c_r^{(j)}|\phi_r\rangle$. Recall from Theorem 15 that for an n -qubit state vector with k terms, the runtime of FASTNORM is $\mathcal{O}(kn^3\epsilon_{\text{FN}}^{-2})$. From the previous discussion, if we select the j th pure state in the decomposition of ρ , we set $k \propto \|c^{(j)}\|_1^2\delta^{-1}$ and $\epsilon_{\text{FN}} \propto \delta w^{-1}$. In a single run of the full algorithm, FASTNORM is called $\mathcal{O}(w)$ times. Therefore, the runtime to generate a single w -length bit string is $T = \mathcal{O}(\|c^{(j)}\|_1^2 w^3 n^3 \delta^{-3})$ with probability p_j . So from Eq. (123), the *average-case* runtime is $\mathcal{O}(\tilde{\Xi} w^3 n^3 \delta^{-3})$. Through $\tilde{\Xi}$, this average-case runtime is sensitive to the particular decomposition of ρ supplied to the simulator. In the case where the decomposition is optimal with

respect to the mixed-state extent Ξ (Definition 3), we have $\tilde{\Xi} = \Xi(\rho)$, so that the average-case runtime is linear in $\Xi(\rho)$. Recall from Sec. III that all single-qubit states admit an equimagical decomposition (Theorem 3) that naturally extends to all tensor products of single-qubit states. In that case $\|c^{(j)}\|_1^2 = \Xi(\rho)$ for all j , so that we can give the *worst-case* runtime as $\mathcal{O}[\Xi(\rho)]$.

The runtime scaling of $\mathcal{O}(\delta^{-3})$ holds provided that the sparsification error δ_S is not smaller than the critical threshold $\delta_c = 8(C - 1)/\|c\|_1^2$, where C is defined in Eq. (104). However, the algorithm is still valid for the case of arbitrary precision, $\delta_S < \delta_c$. In this case we recover the same leading-order scaling as Bravyi *et al.*, namely $\mathcal{O}(\delta^{-4})$ [25], but typically with a prefactor improved by several orders of magnitude (see Fig. 6). A detailed technical analysis is provided in Appendix H, including proof of the following theorem, which captures the results discussed above.

Theorem 19. *Let $\rho = \sum_j p_j |\psi_j\rangle\langle\psi_j|$ be an n -qubit state where every pure state has a known stabilizer decomposition $|\psi_j\rangle = \sum_r c_r^{(j)}|\phi_r\rangle$. For every $|\psi_j\rangle$, let $C_j = \|c^{(j)}\|_1 \sum_r |c_r^{(j)}| \|\langle\psi|\phi_r\rangle\|^2$. Let $\tilde{\Xi} = \sum_j p_j \|c^{(j)}\|_1^2$, and let $D = \max\{(C_j - 1)/\|c^{(j)}\|_1^2\}$. Then for any $p_{\text{fail}} > 0$, and $\delta \geq 24D$ there exists a classical algorithm that, with success probability $(1 - p_{\text{fail}})$, samples a bit string \vec{x} of length w with probability $P_{\text{sim}}(\vec{x})$ such that*

$$\|P_{\text{sim}} - P\|_1 \leq \delta + \mathcal{O}(\delta^2), \quad (124)$$

where $P(\vec{x}) = \text{Tr}(\Pi_{\vec{x}}\rho)$, and $\Pi_{\vec{x}} = |\vec{x}\rangle\langle\vec{x}| \otimes \mathbb{1}_{n-w}$ is a projector. The algorithm returns \vec{x} with random runtime T where the average runtime is

$$\mathbb{E}(T) = \mathcal{O}[w^3 n^3 \tilde{\Xi} \delta^{-3} \log(w/p_{\text{fail}})]. \quad (125)$$

If the decomposition of ρ is optimal with respect to Definition 9, then the expected runtime is $\mathcal{O}[\Xi(\rho)]$. Moreover, if the state decomposition is equimagical, then the right side of Eq. (125) also bounds the worst-case runtime.

If arbitrary precision $\delta \leq 24D$ is required, this can be achieved at the cost of an increased runtime:

$$\mathbb{E}(T) = \mathcal{O}[w^3 n^3 \tilde{\Xi} (\delta^{-3} + 3D\delta^{-4}) \log(w/p_{\text{fail}})]. \quad (126)$$

C. Constrained path simulator

In a standard quasiprobability simulator, the target state $\rho = \sum_j q_j \sigma_j$ is decomposed as an affine combination of frame elements σ_j that are in some sense easy to simulate. We can alternatively combine all the positive and negative contributions into convex combinations σ_+ and σ_- , respectively, so that the decomposition is rewritten: $\rho = \lambda\sigma_+ - (\lambda - 1)\sigma_-$, for some $\lambda = \sum_{q_j \geq 0} q_j \geq 1$. The standard sampling procedure for estimating $\langle E \rangle$ for some

observable E can then be divided into two steps: (i) randomly sample the positive or negative path with probability $\lambda/\|q\|_1$ or $(\lambda - 1)/\|q\|_1$, where $\|q\|_1 = 2\lambda - 1$; (ii) sample an individual frame element from the selected convex combination σ_{\pm} . As explained in Sec. VII A1, the number of samples needed to achieve any accuracy $\epsilon > 0$ is $\mathcal{O}(\|q\|_1^2 \epsilon^{-2})$. Viewed in this way, we see that any increased runtime for simulating magic states arises in step (i) rather than step (ii). In other words, sampling a frame element from the convex combination σ_{\pm} does not incur additional overhead.

An alternative strategy is to constrain sampling to the positive path so that step (i) is avoided. This is equivalent to making the approximation $\rho \approx \lambda\sigma_+$, and comes at the cost of an unavoidable systematic error of size $|(\lambda - 1)\text{Tr}[E\sigma_-]|$. However, an advantage to this approach is that since $\text{Tr}[E\sigma_-]$ is no longer evaluated explicitly, σ_- need not be an efficiently simulable state. Therefore, it is natural to connect this strategy with primal solutions to the generalized robustness problem:

$$\rho = \Lambda^+(\rho)\sigma - [\Lambda^+(\rho) - 1]\rho_-, \quad (127)$$

where σ is a mixed stabilizer state, but ρ_- can be any density operator. Moreover, since systematic error is unavoidable, it is unnecessary to evaluate the first term to high precision, so the runtime can be reduced. Pseudocode for our constrained path simulator is given in Algorithm 1. We then place tight bounds on the systematic error.

Choosing E_{\max} and E_{\min} to be given in steps 3 and 4, we ensure that for all λ and E_{σ} ,

$$|\widehat{E} - \text{Tr}(EO[\rho])| \leq \Delta \quad (128)$$

holds with probability $1 - p_{\text{fail}}$. The major caveat is that there are certain regimes (for large λ and small E_{σ}) where the algorithm fails by trivially estimating the true expectation value to be anywhere in the range $[-1, 1]$. Nevertheless, in some regimes we efficiently obtain a biased but nontrivial estimate. We first briefly explain steps 3 and 4, before analysing the error bound and runtime.

Input: Target state ρ ; real numbers $\lambda, c, p_{\text{fail}} > 0$, and stabilizer state $\sigma \in \overline{\mathcal{S}}_n$ s.t. $\rho \leq \lambda\sigma$ and $c, p_{\text{fail}} \ll 1$; Pauli observable E and simulable channel $O \in \mathcal{O}_n$.

Output: Estimate \widehat{E} and error bound Δ , s.t. $|\widehat{E} - \text{Tr}(EO[\rho])| \leq \Delta$.

-
- 1: $\epsilon \leftarrow c\lambda$
 - 2: Let E_{σ} be an estimate for $\lambda\text{Tr}(EO[\sigma])$ obtained via sampling up to error ϵ and success probability $1 - p_{\text{fail}}$.
 - 3: $E_{\max} \leftarrow \min\{1, E_{\sigma} + \epsilon + \lambda - 1\}$.
 - 4: $E_{\min} \leftarrow \max\{-1, E_{\sigma} - \epsilon - \lambda + 1\}$.
 - 5: $\widehat{E} \leftarrow (E_{\max} + E_{\min})/2$
 - 6: $\Delta \leftarrow (E_{\max} - E_{\min})/2$
-

Algorithm 1. Constrained path simulator.

When λ and σ are such that $\rho \leq \lambda\sigma$, using Eq. (13), there is some density matrix ρ_- such that ρ can be written as

$$\rho = \lambda\sigma - (\lambda - 1)\rho_-. \quad (129)$$

Step 2 estimates E_{σ} such that $|E_{\sigma} - \lambda\text{Tr}(EO[\sigma])| \leq \epsilon$ with probability $1 - p_{\text{fail}}$. We use this to bound possible values of $\text{Tr}(EO[\rho])$:

$$\text{Tr}(EO[\rho]) = \lambda\text{Tr}(EO[\sigma]) - (\lambda - 1)\text{Tr}(EO[\rho_-]) \quad (130)$$

$$\leq E_{\sigma} + \epsilon + (\lambda - 1). \quad (131)$$

Similarly one obtains $\text{Tr}(EO[\rho]) \geq E_{\sigma} - \epsilon - (\lambda - 1)$. Trivially we know that $|\text{Tr}(EO[\rho])| \leq 1$, so in case either expression exceeds this (for example, if E_{σ} is close to ± 1) we simply take either $E_{\max} = 1$ or $E_{\min} = -1$ as necessary. We now consider the regimes where the bounds are trivial, and give the size of the error otherwise.

Case 1 (failure): Trivial bounds are obtained when both these conditions hold:

$$E_{\sigma} + \epsilon + (\lambda - 1) \geq 1, \quad (132)$$

$$E_{\sigma} - \epsilon - (\lambda - 1) \leq -1, \quad (133)$$

that is, when E_{σ} satisfies

$$2 - \lambda(1 + c) \leq E_{\sigma} \leq \lambda(1 + c) - 2. \quad (134)$$

This holds only if $\lambda \geq 2/(1 + c) \approx 2$, as otherwise at most one of the inequalities, Eqs. (132) and (133), can be true.

Case 2 (constant error): When $\lambda < 2/(1 + c) \approx 2$, there is a range of values of E_{σ} where inequalities, Eqs. (132) and (133), are both violated:

$$\lambda(1 + c) - 2 \leq E_{\sigma} \leq 2 - \lambda(1 + c). \quad (135)$$

In this case, we have

$$\widehat{E} = E_{\sigma}, \quad (136)$$

$$\Delta = \lambda(1 + c) - 1. \quad (137)$$

Case 3 (error decreases with $|E_{\sigma}|$): The remaining case occurs when $|E_{\sigma}|$ is sufficiently large, so that either $E_{\max} = 1$ or $E_{\min} = -1$. This limits the range of possible values of $\text{Tr}(EO[\rho])$, so that

$$\Delta = \frac{\lambda(1 + c) - |E_{\sigma}|}{2}. \quad (138)$$

This occurs when exactly one of the inequalities, Eqs. (132) and (133), is satisfied, while the other is violated. Note that this can happen even when $\lambda \gg 2$, as it depends on the value of E_{σ} returned. For example, if $E_{\sigma} = \pm\lambda$, we obtain $\widehat{E} = 1 \mp \epsilon/2$ and $\Delta = \epsilon/2$.

Estimating $\lambda \text{Tr}(EO[\sigma])$ using any Clifford simulator (e.g., the dyadic frame simulator) takes up the most time in the algorithm, as the other steps are trivial to evaluate. Since σ is a convex combination of stabilizer states, there is no additional sampling overhead due to negativity. The prefactor λ increases the variance of the estimator, but we compensate for this by setting the precision to $\epsilon = c\lambda$, where c is a small constant. The rationale for this is that the systematic error due to our ignorance of ρ_- is unavoidable, and this error is of size $\lambda - 1$. Therefore, there is a limit to the precision we can achieve by increasing the runtime of the sampling step, and we should set the precision commensurate with the size of λ . Using the standard arguments (see Sec. VII A1), the smallest number of samples T sufficient to achieve this precision is

$$T = \lceil 2\lambda^2 \epsilon^{-2} \log(2p_{\text{fail}}^{-1}) \rceil = \lceil 2c^{-2} \log(2p_{\text{fail}}^{-1}) \rceil. \quad (139)$$

The runtime for our constrained path simulator is therefore constant with respect to λ [i.e., the generalized robustness $\Lambda^+(\rho)$ when the decomposition is optimal], depending only on the parameters c and p_{fail} . In this sense, we achieve efficient runtime by trading off against precision in the estimate; it is the error Δ that scales with the magic monotone rather than the runtime.

Our explicit algorithm for estimating Pauli expectation values easily adapts to estimate Born-rule probabilities for stabilizer projectors Π by replacing the assumption $|\text{Tr}(E\rho)| \leq 1$ for any ρ with $0 \leq \text{Tr}(\Pi\rho) \leq 1$.

VIII. APPLICATIONS TO OTHER RESOURCES

Although we focus on the simulation of quantum circuits within the stabilizer formalism in Sec. VII, our methods can be extended beyond magic-state quantum computation. The crucial idea here is to identify an efficiently simulable quantum subtheory, consisting of a set of n -qubit pure-state vectors S_n and a set of operators \mathcal{T}_n such that any $K \in \mathcal{T}_n$ acts on a state $|\phi\rangle \in S_n$ as $K|\phi\rangle \propto |\phi'\rangle \in S_n$, and this update can be efficiently tracked. Any such subtheory can then be extended to mixed states and to the dyadic setting, allowing for the adaptation of our classical simulators. In particular, we can show that—as long as the subtheory itself satisfies some basic criteria regarding its simulability—we can always efficiently simulate quantum circuits built from operators in \mathcal{T}_n when acting on states composed of convex mixtures of projectors in S_n . Just as before, when the algorithms work outside the given quantum subtheory, they incur an additional resource cost, which can be measured using our monotones.

This formalism very naturally fits into the framework of quantum resource theories [50], which study the quantification and manipulation of resources in physically restricted settings. Here, a set of states and a set of operations are considered “free,” while states and operations

outside of these sets are costly to use and implement. The connection we build between classical simulators and resource theories then connects the quantitative value of such resources with the performance of the classical simulators, thus giving an explicit operational meaning to important resource monotones. Indeed, the monotones \mathcal{R} , Λ , and Λ^+ can be defined in general resource theories [55], but their operational meaning is not always known. For instance, although the robustness monotones \mathcal{R} and Λ^+ have found general use in tasks such as channel discrimination [75,76] and resource conversion [66,77], the dyadic negativity Λ has not been shown to have any direct operational applications in general resource theories, nor has a connection between monotones such as the generalized robustness Λ^+ and classical simulation been established.

In the following, we refer to the pure states S_n as free. Similarly, we define the set of free operations \mathcal{O}_n as all quantum channels whose Kraus operators belong to \mathcal{T}_n , and thus cannot generate any resource from a free state. Analogously, the set of free observables \mathcal{M}_n can be defined to be all observables, which always result in a free postmeasurement state.

From our discussion of the dyadic frame simulator in Sec. VII A and Appendix C, the proof clearly requires only three crucial assumptions about the classical simulability of the underlying subtheory. We formalize them as follows.

(S1) Only $\mathcal{O}[\text{poly}(n)]$ bits of information are necessary to index all n -qubit pure free states in the set S_n .

(S2) Given a free operator $K \in \mathcal{T}_n$ and any free state $|\phi\rangle \in S_n$, we can compute the update $K|\phi\rangle$ as well as the norm $\|K|\phi\rangle\|$ in $\mathcal{O}[\text{poly}(n)]$ time.

(S3) Given a free observable $\Pi \in \mathcal{M}_n$ and any free states $|L\rangle, |R\rangle \in S_n$, we can compute $\langle R|\Pi|L\rangle$ in $\mathcal{O}[\text{poly}(n)]$ time.

As before, we are interested in the composition of free operations $O \in \mathcal{O}_n$, which admit a simulable decomposition, i.e., can be written as $O(\cdot) = \sum_{i=1}^{N_K} K_i \cdot K_i^\dagger$ with each $K_i \in \mathcal{T}_n$ and $N_K \leq \text{poly}(n)$. With this, our proofs of Sec. VII A and Appendix C can be immediately applied to generalize the dyadic frame simulator.

Theorem 14[†]. *Consider a resource theory with free pure states S_n , free operations \mathcal{O}_n , and free observables \mathcal{M}_n satisfying criteria (S1)–(S3) above.*

Let $\rho = \sum_j \alpha_j |L_j\rangle\langle R_j|$, be a known dyadic decomposition of an initial n -qubit state, where $\alpha \in \mathbb{C}$ and the probability distribution $\{|\alpha_j|/\|\alpha\|_1\}$ can be efficiently sampled. Let $\mathcal{E} = O^{(T)} \circ \dots \circ O^{(1)}$, where each $O^{(i)} \in \mathcal{O}_n$. Suppose that every $O^{(i)}$ has a known simulable decomposition. Then, given $\Pi \in \mathcal{M}_n$, we can estimate the Born-rule probability $\mu = \text{Tr}(\Pi\mathcal{E}[\rho])$ within additive error ϵ ,

with success probability at least $1 - p_{\text{fail}}$ and worst-case runtime

$$\frac{\|\alpha\|_1^2}{\epsilon^2} \log(p_{\text{fail}}^{-1}) T_{\text{poly}}(n). \quad (140)$$

Furthermore, if the dyadic decomposition of ρ is optimal then $\|\alpha\|_1$ can be replaced by $\Lambda(\rho)$.

Theorem 14' establishes an efficient simulation algorithm, which can be employed in any resource theory that satisfies the requirements. The theorem also connects the monotone Λ with the sampling overhead of the algorithm, thus endowing Λ with an exact operational interpretation in the context of resource theories beyond magic. The result of Theorem 14' additionally allows us to employ the constrained path simulator of Sec. VII C to define a related simulation algorithm, which depends on another monotone—the generalized robustness Λ^+ . Once again, the reasoning of Sec. VII C can be applied verbatim under the assumptions (S1)–(S3).

As an example where the result can be immediately applied, consider the resource theory of quantum coherence [51,52], where the free states S_n are the vectors of the computational basis $\{|i\rangle\}$. The free measurements are in the computational basis, for instance projectors of the form $\Pi = |\mathbf{x}\rangle\langle\mathbf{x}| \otimes \mathbb{1}$ where \mathbf{x} is a fixed bit string, which can be efficiently computed. The corresponding dyadic negativity Λ is then the (elementwise) ℓ_1 norm, $\|\rho\|_{\ell_1} = \sum_{i,j} |\langle i|\rho|j\rangle|$. We remark that $\|\cdot\|_{\ell_1}$ is trivially a multiplicative monotone in any dimension. Although $\|\cdot\|_{\ell_1}$ is one of the most commonly employed measures in the resource theory of coherence [51,52,78], it has lacked an explicit operational interpretation thus far. Since the resource theory of coherence is not known to admit a unique, physically motivated choice of free operations [52,79], we briefly discuss the possible choices of \mathcal{O}_n and their classical simulability. From this, we use Theorem 14' to give $\|\cdot\|_{\ell_1}$ an operational interpretation.

The most fundamental class of free operations within the resource theory of coherence are the *incoherent operations* (IOs) [51], defined to be maps that admit a decomposition into Kraus operators that preserve the set of incoherent states. Such Kraus operators can be expressed as [80] $K = \sum_{x \in S} c_x |f(x)\rangle\langle x|$ for some set of bit strings S , coefficients c_x , and an arbitrary function f . Given such a K acting on no more than b qubits, where b is constant, we can efficiently compute $(K \otimes \mathbb{1}_{n-b})|i\rangle$. Since any Boolean f can be implemented by composing a set of universal classical logic gates (with $b = 2$) such as AND and XOR, such gates can generate an IO realizing any Boolean function. Furthermore, IOs can simulate any quantum channel with sufficiently many coherent states [51,81].

One family of useful IOs in practice are the *strictly incoherent operations* (SIOs) [80,82], which can be efficiently

implemented by quantum circuits using only incoherent ancillae [82]. As a subtheory of IO, all the updates are still efficiently computable. Furthermore, $b = 3$ suffices to provide the Toffoli gate, which is universal for classical reversible logic, and so can generate any Kraus operator of the required form. The biggest difference between SIOs and IOs is that, while SIOs are better understood from the perspective of their practical implementation, they cannot be promoted to universal quantum operations through the use of ancillary resource states [81].

We conclude that the resource theory of coherence that uses either IOs or SIOs as free operations satisfies the conditions of Theorem 14'. Thus, the theorem endows the ℓ_1 norm of coherence with an operational interpretation as the sampling overhead in black the simulation of either of the classes of operations using the dyadic frame simulator, and similarly the constrained path simulator gives another meaning to the robustness of coherence Λ^+ [83].

Our dyadic frame simulator is especially useful in resource theories where other simulation algorithms such as the Howard-Campbell simulator for magic states [21] cannot be readily adapted. For instance, in the resource theory of coherence, the free states S_n form a zero-measure subset of all states, which means that no resourceful state ρ can be decomposed as $\rho = \sum_j p_j |\phi_j\rangle\langle\phi_j|$ with $|\phi_j\rangle \in S_n$ and so the corresponding robustness quantifier $\mathcal{R}(\rho)$ diverges.

Note, however, that the dyadic frame simulator does not work for all resource theories. While the dyadic frames for stabilizer and incoherent operations meet the conditions of Theorem 14, the requirements cannot hold for the theory of separable states under local operations and classical communication (LOCC). This is because the free states consist of an infinite number of inequivalent pure product states, which cannot be described using $\text{poly}(n)$ bits. However, one can accurately compute local unitaries acting on product states efficiently. Indeed, our framework could encompass entanglement and similar theories using a suitable ϵ net over the set of separable states, and we leave the precise statement of the relevant conditions for future work.

IX. CONCLUSIONS

We introduce three resource monotones into the setting of magic state quantum computation: the dyadic negativity Λ , the generalized robustness Λ^+ , and the mixed-state extent Ξ . The first part of the paper focuses on resource-theoretic results, including that (i) for pure states, the monotones all equal the extent monotone ξ ; (ii) for tensor products of single-qubit mixed states, they all coincide; and (iii) the monotones act multiplicatively on tensor products of single-qubit mixed states. The results significantly simplify the computation of the monotones for multiple copies of a single-qubit state, and allow us to completely

understand the asymptotic behavior of our magic quantifiers, which contrasts with previously used monotones. Furthermore, our magic monotones often tighten previously known bounds on distillation rates.

For each monotone, we introduce a related classical simulation algorithm. Our dyadic negativity simulator has a runtime proportional to $\Lambda(\rho)^2$, which is similar to—but significantly faster than—the Howard-Campbell simulator with runtime $\mathcal{R}(\rho)^2$ where \mathcal{R} is the robustness of magic. Additionally, we show that the dyadic negativity simulator works for circuits that use completely stabilizer-preserving operations. This class includes all the conventional stabilizer operations (Clifford unitaries, Pauli projections, etc.) and we believe it is likely to be strictly larger. If true, the situation would mirror entanglement theory in the separation between LOCC and separable operations [84]. We find that for tensor products of n single-qubit states, both Λ and \mathcal{R} scale exponentially with n , but \mathcal{R} is exponentially larger than Λ . This establishes our dyadic negativity simulator as the fastest known quasiprobability simulator for qubit magic states.

However, not all classical simulation algorithms are based on quasiprobability distributions, with the stabilizer-rank methods representing a distinct paradigm. There are several crucial differences, including that stabilizer-rank methods enable a stronger notion of classical simulation, as they allow us to sample outputs of a quantum computation, not just estimate Born-rule probabilities. Prior work on stabilizer-rank simulations considered only pure states, but our simulator extends this to mixed states and demonstrates an expected runtime proportional to $\Xi(\rho)$. Note the linear dependence on Ξ (largely due to fast norm estimation [24]), in contrast to the quadratic dependence encountered with quasiprobability simulators. Since in general $\Lambda[\rho] \leq \Xi[\rho]$, it is theoretically possible that $\Lambda[\rho] \ll \Xi[\rho]$ so that $\Lambda[\rho]^2 \leq \Xi[\rho]$, which would mean a runtime advantage for the quasiprobability methods. However, for products of single-qubit states the monotones are equal, so for such states our resource theory results show that the advantage clearly falls to the stabilizer-rank simulators. Furthermore, we improve stabilizer-rank bounds, with the runtime for sampling Clifford magic states (e.g., T states) improved to $\mathcal{O}(1/\delta^3)$ from the prior $\mathcal{O}(1/\delta^4)$ bound where δ is the sampling precision. For other magic states, the advantage is not as simple to describe using big- \mathcal{O} notation, but Fig. 6 shows it to be considerable in practice.

Finally, by ensuring that our simulation algorithms can be easily generalized and providing a recipe to adapt the simulators to resource theories beyond magic states, we shed light on the simulation of quantum circuits using very general resources under suitable assumptions. This not only provides new insight into the practical uses of resource quantifiers in well-studied theories such as quantum coherence, but also opens an avenue for a further study of the connections between the theoretical frameworks of

quantum resources and their operational applications in quantum computation.

A clear direction for further research is to extend our results to the channel picture, which would enable a more direct route to simulate circuits with no need to replace nonfree operations with state injection gadgets. This is especially important in the context of stabilizer theory for the simulation of circuits with gates outside the Clifford hierarchy, as the gadgets then become more complex.

ACKNOWLEDGMENTS

We would like to thank the anonymous reviewers for their insightful comments. This work is supported by the Engineering and Physical Sciences Research Council [Grants No. EP/P510270/1 (J.R.S.) and No. EP/M024261/1 (E.T.C. and Y.O.)]. B.R. is supported by the Presidential Postdoctoral Fellowship from Nanyang Technological University, Singapore. Research at Perimeter Institute is supported in part by the Government of Canada through the Department of Innovation, Science and Economic Development Canada and by the Province of Ontario through the Ministry of Colleges and Universities. H.P. also acknowledges the support of the Natural Sciences and Engineering Research Council of Canada (NSERC) discovery Grants [RGPIN-2019-04198] and [RGPIN-2018-05188] and the ARC via the Centre of Excellence in Engineered Quantum Systems (EQUS) Project No. CE170100009. This work was completed while E.T.C. was at the University of Sheffield.

APPENDIX A: ALTERNATIVE PROOFS FOR PREVIOUS RESULTS

1. Monotone equivalence proof

Here we prove Lemma 1. Consider an optimal decomposition for the extent, such that

$$|\Psi\rangle = \sum_i c_i |\psi_i\rangle, |\psi_i\rangle \in S_n, \quad (\text{A1})$$

with $\xi(\Psi) = \|c\|_1^2$. Then

$$|\Psi\rangle\langle\Psi| = \sum_{ij} c_i c_j^* |\psi_i\rangle\langle\psi_j| \quad (\text{A2})$$

is a valid decomposition into the dyadic frame leading to

$$\begin{aligned} \Lambda(|\Psi\rangle\langle\Psi|) &\leq \sum_{ij} |c_i c_j| \quad (\text{A3}) \\ &= \left(\sum_i |c_i| \right) \left(\sum_j |c_j| \right) \\ &= \|c\|_1^2 = \xi(\Psi). \end{aligned}$$

Next, we prove the converse inequality. The dual convex problem to the minimization of ξ is

$$\xi(\Psi) = \max_{\omega} \{ |\langle \omega | \Psi \rangle|^2 : \text{such that } \forall \phi \in \mathcal{S}_n, |\langle \omega | \phi \rangle|^2 \leq 1 \}. \quad (\text{A4})$$

Note that the ω need not be properly normalized. Let us label ω^* as a vector achieving this maximum so that $\xi(\Psi) = |\langle \omega^* | \Psi \rangle|^2$. We further recall that Λ also has a dual formulation

$$\begin{aligned} \Lambda(\rho) &= \max_W \{ \text{Tr}[W\rho] : \text{such that } \forall |\phi\rangle, |\psi\rangle \\ &\in \mathcal{S}_n, |\langle \phi | W | \psi \rangle| \leq 1 \}. \end{aligned} \quad (\text{A5})$$

In particular for feasible W we have $\Lambda(\rho) \geq \text{Tr}[W\rho]$. We notice that the extent witness $|\omega^*\rangle$ can be used to build an operator $W = |\omega^*\rangle\langle\omega^*|$ that is a valid witness for Λ . Therefore,

$$\begin{aligned} \Lambda(|\Psi\rangle\langle\Psi|) &\geq \text{Tr}[|\omega^*\rangle\langle\omega^*| |\Psi\rangle\langle\Psi|] \\ &\geq |\langle \omega^* | \Psi \rangle|^2 = \xi(\Psi). \end{aligned} \quad (\text{A6})$$

Having proved both directions, we conclude an equality. This proves Lemma 1. Since the witness W is a positive operator, an identical proof also shows that $\Lambda^+(|\Psi\rangle\langle\Psi|) = \xi(|\Psi\rangle\langle\Psi|)$. For the Ξ monotone, there is only one convex decomposition of $|\Psi\rangle\langle\Psi|$. Hence $\Xi(|\Psi\rangle\langle\Psi|) = \xi(|\Psi\rangle\langle\Psi|)$.

2. Sandwich theorem

Here we present a proof of Theorem 2. Recall that Λ is the result of maximizing over all W witnesses, whereas Λ^+ is limited to all W^+ witnesses, which immediately leads to $\Lambda^+(\rho) \leq \Lambda(\rho)$. To show $\Lambda(\rho) \leq \Xi(\rho)$, one simply takes the optimal decomposition with respect to Ξ , as follows:

$$\begin{aligned} \rho &= \sum_j p_j |\Psi_j\rangle\langle\Psi_j|, \\ \Xi[\rho] &= \sum_j p_j \xi(\Psi_j). \end{aligned} \quad (\text{A7})$$

Next, we insert this decomposition into Λ and use convexity

$$\Lambda[\rho] \leq \sum_j p_j \Lambda(|\Psi_j\rangle\langle\Psi_j|). \quad (\text{A8})$$

Using Lemma 1 we have

$$\Lambda[\rho] \leq \sum_j p_j \xi(\Psi_j) = \Xi[\rho], \quad (\text{A9})$$

which completes the proof of Theorem 2.

APPENDIX B: GEOMETRY OF Φ_ρ^\pm

Here we show that the states Φ_ρ^\pm introduced in Eq. (49) are contained in the set P_Y introduced in Definition 7. Note that Eq. (49) is defined in terms of $\rho \in P_Y$.

The result $\Phi_\rho^\pm \in P_Y$ is used in the proof of Lemma 7. In that proof, we appeal to geometry presented in Fig. 3 and here we instead provide an algebraic argument.

Using the definition of the set P_Y , to prove that $\Phi_\rho^\pm \in P_Y$, it suffices to show that

$$\begin{aligned} \langle \Phi_\rho^\pm | Y | \Phi_\rho^\pm \rangle &\geq 0, \\ \langle \Phi_\rho^\pm | (X - Y) | \Phi_\rho^\pm \rangle &\geq 0, \\ \langle \Phi_\rho^\pm | (Z - Y) | \Phi_\rho^\pm \rangle &\geq 0. \end{aligned} \quad (\text{B1})$$

Now note that $\langle \Phi_\rho^\pm | Y | \Phi_\rho^\pm \rangle = \text{Tr}[Y\rho]$. Minimizing $\text{Tr}[Y\rho]$ over all feasible Bloch vectors (r_A, r_B, r_F) in the decomposition of ρ given by Eq. (46), we find that $\text{Tr}[Y\rho] \geq 0$, which proves the first inequality.

Next, we tackle the second inequality (with the third inequality following in a similar fashion). One computes that

$$\begin{aligned} 2\langle \Phi_\rho^\pm | (X - Y) | \Phi_\rho^\pm \rangle &= \langle \Phi_\rho^\pm | (\sqrt{6}\sigma_A + \sqrt{2}\sigma_B) | \Phi_\rho^\pm \rangle \\ &= \sqrt{6}r_A \pm \sqrt{2}\sqrt{1 - r_A^2 - f^2}, \end{aligned} \quad (\text{B2})$$

which is positive whenever

$$\sqrt{6}r_A \geq \sqrt{2(1 - r_A^2 - f^2)} \quad (\text{B3})$$

or more concisely

$$r_A \geq \sqrt{(1 - f^2)}/2. \quad (\text{B4})$$

For mixtures of Ψ_f^X and Ψ_f^Z , we find this holds with equality. For mixtures in the convex hull of $\{\Psi_f^X, \Psi_f^Y, \Psi_f^Z\}$ we find $r_A \leq \sqrt{(1 - f^2)}/2$. However, we are currently considering ρ outside this set, just outside the facet spanned by Ψ_f^X and Ψ_f^Z . Therefore, Eq. (B4) indeed holds.

APPENDIX C: DYADIC FRAME SIMULATOR TECHNICAL DETAILS

Here we prove Theorem 14, which assures the validity and runtime of our dyadic frame simulator. Recall that the goal of the dyadic frame simulator is to estimate the Born-rule probability $\mu = \text{Tr}[\Pi\mathcal{E}(\rho)]$, where ρ is an n -qubit mixed magic state, and $\mathcal{E} = O^{(T)} \circ \dots \circ O^{(1)}$ is a sequence of completely stabilizer-preserving channels $O^{(t)} \in \mathcal{O}_n$. Before proving the theorem, we first restate

and discuss the restrictions we impose on $O^{(t)}$. Any completely stabilizer-preserving channel $O \in \mathcal{O}_n$ has at least one Kraus decomposition of the form

$$O = \sum_r^{N_U} p_r \mathcal{U}_r + \sum_s^{N_K} q_s \mathcal{K}_s, \quad (C1)$$

where all $\mathcal{U}_r = U_r(\cdot)U_r^\dagger$ are unitary Clifford operations, and $\mathcal{K}_s = K_s(\cdot)K_s^\dagger$ correspond to completely stabilizer-preserving nonunitary Kraus operators [22]. Let $P_U = \sum_r p_r$ be the total weight of the unitary part of the decomposition. We say that a channel decomposition is *simulable* if the number of Kraus operators is bounded as $N_U, N_K \leq \text{poly}(n)$. In Theorem 14 we assume that the channels $O^{(t)}$ provided as input to the algorithm all have a known simulable decomposition. We use the restriction on N_U for simplicity in the proof, but provided one can efficiently sample from the distribution $\{p_r/P_U\}$ and compute any corresponding U_r , then this restriction can be removed. Note however that the restriction on N_K cannot be similarly relaxed. For concreteness, we assume that nonunitary Kraus operators are given as a length $\text{poly}(n)$ list with each entry being a pair. The pair encodes a stabilizer-preserving Kraus operator and its associated weight factor. The Clifford part of the decomposition takes the same format. We use $\mathcal{L}_U^{(t)}$ and $\mathcal{L}_K^{(t)}$ to denote the respective lists for the unitary and nonunitary part of the decomposition of $O^{(t)}$. Each stabilizer-preserving Kraus operator K is described by giving an efficient description of the stabilizer state corresponding to the Choi state Φ_K as

$$\Phi_K = (\mathcal{K} \otimes \mathbb{1})|\Phi\rangle\langle\Phi|, \quad (C2)$$

where $|\Phi\rangle \propto \sum_j |j\rangle|j\rangle$. Note that since K acts by conjugation, Φ_K is a pure stabilizer state, so can be specified by $\mathcal{O}(n^2)$ classical bits.

The class of simulable channel decompositions encompasses a wide range of practically important stabilizer operations. First, any convex combination of n -qubit Clifford gates $O = \sum_j p_j U_j$ is included, provided $\{p_j\}$ can be efficiently sampled from. Another subset of simulable channels are those of the form $O = O' \otimes \mathbb{1}_{n-b}$, where $O' \in \mathcal{O}_b$ and b is a small constant. Any $O' \in \mathcal{O}_b$ has a $2b$ -qubit Choi state $\Phi_{O'}$ that lies inside the stabilizer polytope (i.e., it can be written as a convex combination of pure stabilizer states, each corresponding to a Kraus operator) [22]. Although the number of stabilizer states grows superexponentially with b , the real vector space inhabited by $2b$ -qubit density matrices is $(4^b - 1)$ dimensional. We can therefore completely partition the stabilizer polytope into simplices with 4^b vertices, where any mixed stabilizer state inhabits at least one simplex. Hence, by Carathéodory's theorem,

Φ_O can be written as a convex combination of at most 4^b pure stabilizer states. Thus for families of circuits where b has a fixed upper bound, $N_U + N_K$ does not grow with n . This restriction is not too onerous, since practical quantum algorithms are typically synthesized in terms of one-, two-, and three-qubit gates, and noise channels are often assumed to act locally. Moreover, we often already know the stabilizer decompositions of interesting channels. For instance, we can express T -gate injection gadgets and the single-qubit depolarizing channel with only two and four Kraus operators, respectively.

We now present the algorithm and prove its validity. Our algorithm has two subroutines: (i) Algorithm 2, which is an extended Gottesman-Knill-type subroutine that probabilistically updates an input stabilizer dyad given a set of Kraus operators; and (ii) Algorithm 3, which is an outer quasiprobability sampling routine that samples an initial dyad from the initial nonstabilizer state and propagates the dyad through the circuit, randomly selecting a single Kraus operator from each decomposition $O^{(t)}$.

Note that in Algorithm 2, we use the trace norm (i.e., the Schatten 1-norm, $\|A\|_1 = \text{Tr}[\sqrt{A^\dagger A}]$), rather than the usual trace (as in the Born rule) to calculate the transition probabilities for propagating with a particular Kraus operator. While $\text{Tr}(\Pi\rho) = \|\Pi\rho\Pi^\dagger\|_1$ for physical states ρ , this does not hold for general dyads $|L\rangle\langle R|$. We illustrate that the Schatten 1-norm is the appropriate choice with a toy example. Consider the scenario where the penultimate dyad is $\sigma^{(T-1)} = |+\rangle\langle -|$, the final stabilizer channel $O^{(T)}$ is defined by Kraus operators $K_1 = \mathbb{1} \otimes |0\rangle\langle 0|$ and $K_2 = U \otimes |1\rangle\langle 1|$ for some Clifford U , and the final measurement operator to be evaluated is $\Pi = |1\rangle\langle 1| \otimes \mathbb{1}$. Now, the channel $O^{(T)}$ leaves $\sigma^{(T-1)}$ unchanged, $O^{(T)}(|+\rangle\langle -|) = |+\rangle\langle -|$. It is therefore clear that the correct contribution to the expectation value estimate (line 15 in Algorithm 3) should be

$$\mu_m = \|\alpha\|_1 \text{Re}\{\text{Tr}[\Pi\sigma^{(T)}]\}, \quad (C3)$$

$$= \|\alpha\|_1 \text{Re}\{\text{Tr}[(|1\rangle\langle 1| \otimes \mathbb{1})|+\rangle\langle -|]\}, \quad (C4)$$

$$= \|\alpha\|_1 \text{Re}\{\langle -|1\rangle\langle 1|+\rangle\} = -\|\alpha\|_1/2, \quad (C5)$$

where we use cyclicity of the trace, and neglect the phase $e^{i\theta_r}$ for brevity. We need to ensure that the transition probabilities we compute (in line 4 of Algorithm 2) produce statistics that converge to this contribution. Suppose we were to naively use the trace to compute transition probabilities, $P_{\text{Tr},j} = \text{Tr}[K_j\sigma^{(T-1)}K_j^\dagger]$. Then we would obtain

$$P_{\text{Tr},1} = \text{Tr}[(\mathbb{1} \otimes |0\rangle\langle 0|)|+\rangle\langle -|(\mathbb{1} \otimes |0\rangle\langle 0|)], \quad (C6)$$

$$= \langle -|+\rangle = 0, \quad (C7)$$

$$P_{\text{Tr},2} = \text{Tr}[(U \otimes |1\rangle\langle 1|)|+0\rangle\langle -0|(U^\dagger \otimes |1\rangle\langle 1|)], \quad (\text{C8})$$

$$= \langle -|U^\dagger U|+\rangle\langle 1|0\rangle|^2 = 0. \quad (\text{C9})$$

Here we have a problem, because both paths evaluate to zero, preventing μ_m from making any nonzero contribution to our estimate. By contrast, in our algorithm we use the Schatten 1 norm to compute transition probabilities:

$$P_1 = \|(\mathbb{1} \otimes |0\rangle\langle 0|)|+0\rangle\langle -0|(\mathbb{1} \otimes |0\rangle\langle 0|)\|_1, \quad (\text{C10})$$

$$= \| |+0\rangle\langle -0| \|_1 = 1, \quad (\text{C11})$$

$$P_2 = \|(U \otimes |1\rangle\langle 1|)|+0\rangle\langle -0|(U^\dagger \otimes |1\rangle\langle 1|)\|_1, \quad (\text{C12})$$

$$= |\langle 1|0\rangle|^2 \|(|U|+\rangle)|1\rangle\langle 1|((-)|U^\dagger)\|_1 = 0. \quad (\text{C13})$$

This method correctly tells us that we should select Kraus operator K_1 with certainty, resulting in the correct contribution $\mu_m = -\|\alpha\|_1/2$.

Below we prove that this strategy leads to an unbiased estimator for μ , where each individual sample is bounded as $|\mu_m| \leq \|\alpha\|_1$. As per standard quasiprobability simulators (see Sec. VII A1), to estimate an observable within additive error of ϵ with success probability $p_{\text{succ}} \geq 1 - p_{\text{fail}}$, we require at least M samples from our algorithm [20,21], where

$$M \geq 2 \frac{\|\alpha\|_1^2}{\epsilon^2} \log\left(\frac{2}{p_{\text{fail}}}\right). \quad (\text{C14})$$

To prove the validity of our algorithm we must (i) explain how the stabilizer update $q_r K_r \sigma K_r^\dagger$ can be carried out efficiently, (ii) show that the values P_r in steps 2–6 of Algorithm 2 form a proper probability distribution, and (iii) show that $\hat{\mu}$ returned by Algorithm 3 is an unbiased estimator for $\text{Tr}[\Pi \mathcal{E}(\rho)]$. The total runtime given in

Input: Initial stabilizer dyad $\sigma = |L\rangle\langle R|$; length- N list of pairs $\mathcal{L} = \{(q_1, K_1), \dots, (q_N, K_N)\}$, where $q_j > 0$ are weights and K_j are stabilizer Kraus operators; normalization factor P_X .

Output: Updated dyad $\sigma' = |L'\rangle\langle R'|$.

```

1: function STABILIZERUPDATE( $\sigma, \mathcal{L}, P_X$ )
2:   for  $r \leftarrow 1$  to  $N$  do
3:      $q'_r \leftarrow q_r / P_X$ 
4:      $P_r \leftarrow \|q'_r K_r \sigma K_r^\dagger\|_1$ 
5:   end for
6:    $P_0 \leftarrow 1 - \sum_{r=1}^N P_r$ 
7:   Sample  $s$  from  $\{0, \dots, N\}$  with probability  $P_s$ 
8:   if  $s = 0$  then
9:      $\sigma' \leftarrow \sigma$ 
10:  else
11:     $\sigma' \leftarrow |L'\rangle\langle R'| = q'_s K_s |L\rangle\langle R| K_s^\dagger / P_s$ 
12:  end if
13:  return  $\sigma'$ 
14: end function
    
```

Algorithm 2. Stabilizer Kraus update subroutine.

Input: Initial state ρ with known dyadic stabilizer decomposition $\rho = \sum_j \alpha_j |L_j\rangle\langle R_j|$; number of samples M ; list of channels $\{O^{(1)}, \dots, O^{(T)}\}$ of length T where $O^{(t)} \in \mathcal{O}_n$ for all t , and each is *simulable* as described in this appendix, and described by a list of weighted Clifford gates $\mathcal{L}_U^{(t)}$ and non-unitary Kraus operators $\mathcal{L}_K^{(t)}$; stabilizer projector Π .

Output: Estimate $\hat{\mu}$ for $\text{Tr}[\Pi \mathcal{E}(\rho)]$, where $\mathcal{E} = O^{(T)} \circ \dots \circ O^{(1)}$.

```

1: Let  $P_j = |\alpha_j| / \|\alpha\|_1$  define a probability distribution.
2: for  $m \leftarrow 1$  to  $M$  do
3:   Sample  $r$  with probability  $P_r$ .
4:    $\sigma^{(0)} \leftarrow |L^{(0)}\rangle\langle R^{(0)}| = |L_r\rangle\langle R_r|$ ,  $e^{i\theta_r} \leftarrow \alpha_r / |\alpha_r|$ .
5:   for  $t \leftarrow 1$  to  $T$  do
6:      $P_U \leftarrow \sum_j p_j$ , where  $p_j$  are the weights in list  $\mathcal{L}_U^{(t)}$ .
7:      $P_K \leftarrow 1 - P_U$ 
8:     With probability  $P_U$ ,  $X \leftarrow "U"$ , else  $X \leftarrow "K"$ .
9:      $\sigma^{(t)} \leftarrow \text{STABILIZERUPDATE}(\sigma^{(t-1)}, \mathcal{L}_X^{(t)}, P_X)$ 
10:    if  $\sigma^{(t)} = 0$  then
11:       $\sigma^{(T)} \leftarrow 0$ 
12:    break
13:  end if
14: end for
15:  $\mu_m \leftarrow \text{Re}\{\|\alpha\|_1 e^{i\theta_r} \text{Tr}[\Pi \sigma^{(T)}]\}$ .
16: end for
17: return  $\hat{\mu} \leftarrow \sum_m \mu_m / M$ 
    
```

Algorithm 3. Dyadic frame simulator.

Theorem 14 is the product of the number of samples M and the runtime to compute each sample.

(i) *Efficient stabilizer update with Kraus operators.* In Algorithm 2 we must compute the trace norm $\|q_r K_r |L\rangle\langle R| K_r^\dagger\|_1$ for all n pairs $(q_r, K_r) \in \mathcal{L}$, and then perform the update $|L\rangle\langle R| \rightarrow |L'\rangle\langle R'|$. Note that we track any accumulated phase through the update, but here we absorb this factor in $|L'\rangle\langle R'|$ for brevity. There are two cases. Either \mathcal{L} is a list of unitary operators, or it is a list of nonunitary Kraus operators. In the unitary case, computation of the norm is trivial. Since the initial dyad σ is normalized with respect to the trace norm, and the norm is unitarily invariant, we have $\|p_r U_r \sigma U_r^\dagger\|_1 = p_r$. By assumption, $p_r / \sum_r p_r$ can be efficiently sampled from. In the nonunitary case, we must compute

$$\|q'_r K_r |L\rangle\langle R| K_r^\dagger\|_1 = q'_r \|K_r |L\rangle\| \cdot \|K_r |R\rangle\|. \quad (\text{C15})$$

Note that unlike the trace, the trace norm does not depend on the overlap between $K_r |L\rangle$ and $K_r |R\rangle$, and their vector norms are calculated separately. To see how this is done, we note that a Kraus operator whose Choi state is a normalized pure stabilizer state can always be written in the form $K_r = 2^{h/2} U_r \Pi_r$, where U_r is a Clifford gate and Π_r is a stabilizer projector of rank 2^{n-h} [85], for some $h \leq n$. Since U_r leaves the norm invariant, for the purpose of computing transition probabilities, only the projector and

normalization constant matter:

$$\|q'_r K_r |L\rangle\langle R| K_r^\dagger\|_1 = q'_r 2^h \|\Pi_r |L\rangle\| \cdot \|\Pi_r |R\rangle\|. \quad (\text{C16})$$

The projection of each pure state onto a stabilizer subspace can be computed using standard stabilizer simulation techniques [3,25], in time $\mathcal{O}(hn^2)$. In the nonunitary case, we must compute the norm for $2N_K$ projected stabilizer states, so the total runtime for computing all transition probabilities for a single step t is $\mathcal{O}(hN_K n^2)$.

Once all N transition probabilities are computed, a single Kraus operator $K_s \propto U_s \Pi_s$ is randomly selected for the update, so we must compute $|L'\rangle\langle R'| \propto U_s \Pi_s |L\rangle\langle R| \Pi_s U_s^\dagger$. As discussed in Sec. VII A, it is vital that we track any acquired phase throughout each update, as this will affect our final estimate when we average over all M samples μ_m . We can do this using the phase-sensitive Clifford simulator described in Ref. [25]. There it was shown that the update corresponding to the projection $(\mathbb{1} + Q)/2$, where Q is a Pauli operator, can be carried out in $\mathcal{O}(n^2)$ steps. A rank 2^{n-h} stabilizer projector can be decomposed as a product of h Pauli projections, so the projective part of the update takes time $\mathcal{O}(hn^2)$. As for the Clifford update, any n -qubit Clifford operation can be written in canonical form comprised of $\mathcal{O}[n^2/\log(n)]$ gates from the standard gate set {CNOT, H , S } [3]. Bravyi *et al.* [25] showed that their phase-sensitive Clifford simulator can perform CNOT and S updates in time $\mathcal{O}(n)$, and H in time $\mathcal{O}(n^2)$, so the Clifford update for U_r can be completed in time $\mathcal{O}[n^4/\log(n)]$. Since $h \leq n$ and for simulable decompositions $N_K \leq \text{poly}(n)$, the time taken is $\text{poly}(n)$.

Combining all steps, the total time for a single call to STABILIZERUPDATE is $\mathcal{O}[h(N_K + 1)n^2] + \mathcal{O}[n^4/\log(n)]$. Since $h \leq n$ and for simulable decompositions $N_K \leq \text{poly}(n)$, the call is completed in $\text{poly}(n)$ time in the general case. We note that in the special case where we restrict each $O^{(t)}$ to act on at most b qubits, for some fixed b , the runtime for a single call can be improved considerably. In that case, $N_K \leq 4^b$ (i.e., a constant with respect to n) and the runtime will be $\mathcal{O}[b(4^b + 1)n^2] + \mathcal{O}[b^2 n^2/\log(b)]$.

(ii) *Valid probability distribution.* From the definition of P_0 in step 6 of Algorithm 2, it is clear that $\sum_{r=0}^N P_r = 1$ and $P_1, \dots, P_N \geq 0$. Hence, to show that $\{P_r\}$ is a probability distribution, it suffices to show that $P_0 \geq 0$. This is trivially true for the unitary path, as $\left\{p_r / \left(\sum_{s=1}^{N_U} p_s\right)\right\}$ is clearly a properly normalized distribution, with $P_0 = 0$. It remains to show $P_0 \geq 0$ for the nonunitary case.

It is given that the channel $O^{(t)}$ is a completely positive and trace-preserving (CPTP) map. Let $O_K(\cdot) = \sum_{r=1}^{N_K} q_r K_r(\cdot) K_r^\dagger$ denote the nonunitary part of the decomposition, Eq. (C1), and let $O_U(\cdot) = \sum_{r=1}^{N_U} p_r U_r(\cdot) U_r^\dagger$, so that $O^{(t)} = O_U + O_K$. Each Clifford is a unitary operator, so $\sum_{r=1}^{N_U} p_r U_r^\dagger U_r = P_U \mathbb{1}$, recalling that $P_U = \sum_{r=1}^{N_U} p_r$ and $P_K = 1 - P_U$. Since $O^{(t)}$ is CPTP, its Kraus representation

must be complete:

$$\mathbb{1} = P_U \mathbb{1} + \sum_{r=1}^{N_K} (\sqrt{q_r} K_r)^\dagger (\sqrt{q_r} K_r). \quad (\text{C17})$$

It follows that $\sum_{r=1}^{N_K} (\sqrt{q_r} K_r)^\dagger (\sqrt{q_r} K_r) = P_K \mathbb{1}$, meaning that O_K is a CPTP map up to normalization by $1/P_K$. This normalization is achieved by setting $q'_r = q_r/P_K$ in step 3. Then for any pure state $|\psi\rangle$, we have

$$\begin{aligned} 1 &\geq \text{Tr}[O(|\psi\rangle\langle\psi|)/P_K] = \sum_{r=1}^{N_K} \text{Tr}[\sqrt{q'_r} K_r |\psi\rangle\langle\psi| K_r^\dagger \sqrt{q'_r}], \\ &= \sum_{r=1}^{N_K} \|\sqrt{q'_r} K_r |\psi\rangle\|^2. \end{aligned} \quad (\text{C18})$$

Let $Q^{(\psi)}$ be the N_K -element real vector where the r th entry is $Q_r^{(\psi)} = \|\sqrt{q'_r} K_r |\psi\rangle\|$. From Eq. (C18), we have that $\|Q^{(\psi)}\| \leq 1$. Then for any normalized dyad $|L\rangle\langle R|$ we can express the sum of P_r for $r \geq 1$ as a dot product between $Q^{(L)}$ and $Q^{(R)}$:

$$\sum_{r=1}^{N_K} P_r = \sum_{r=1}^{N_K} \|q'_r K_r |L\rangle\langle R| K_r^\dagger\|_1, \quad (\text{C19})$$

$$= \sum_{r=1}^{N_K} \|\sqrt{q_r} K_r |L\rangle\| \cdot \|\sqrt{q_r} K_r |R\rangle\|, \quad (\text{C20})$$

$$= \sum_{r=1}^{N_K} Q_r^{(L)} Q_r^{(R)}, \quad (\text{C21})$$

$$= Q^{(L)} \cdot Q^{(R)} \leq \|Q^{(L)}\| \cdot \|Q^{(R)}\| \leq 1, \quad (\text{C22})$$

where in the last line we use the Cauchy-Schwarz inequality to show that $\sum_{r \geq 1} P_r \leq 1$, as promised. We note that the strategy of using an ‘‘abort’’ outcome P_0 is deployed in the appendix of Ref. [28] to simulate postselective channels. In our case the fact that P_r for $r \geq 1$ can sum to less than 1 instead arises from the non-Hermiticity of the initial dyad σ .

(iii) *Unbiased estimator.* Finally we show that the expected value of $\hat{\mu}$ in Algorithm 3 is $\text{Tr}[\Pi \mathcal{E}(\rho)]$. First, let us recombine the unitary and nonunitary part of $O^{(t)}$ into a single Kraus representation $O^{(t)} = \sum_{a=1}^{N_{\text{Tot}}} q_r^{(t)} K_r^{(t)}(\cdot) K_r^{(t)\dagger}$, where $N_{\text{Tot}} = N_U + N_K$, and consider the probability of sampling the r th pair (q_r, K_r) at step t . By inspection of Algorithms 2 and 3 we see that the probability of taking the unitary path and then selecting the r th pair from \mathcal{L}_U is

$$\text{Pr}(\text{‘‘}U\text{’’}, r) = P_U \times P_r = P_U \|p'_r U_r \sigma U_r^\dagger\|_1 = \|p_r U_r \sigma U_r^\dagger\|_1. \quad (\text{C23})$$

Similarly the probability of choosing the nonunitary path followed by the r th element of \mathcal{L}_K is $\text{Pr}(\text{‘‘}K\text{’’}, r) = P_K P_r =$

$\|q_r K_r \sigma K_r^\dagger\|_1$. Thus at step t , the probability of sampling any Kraus operator from the decomposition, whether unitary or nonunitary, is given by $\Pr(q_r, K_r) = \|q_r K_r \sigma K_r^\dagger\|_1$, and we can drop the distinction between the two.

Now, let the $(T+1)$ -element vector $\vec{r} = (r_0, r_1, \dots, r_T)$ label a particular trajectory through the circuit, in the following sense. The first entry r_0 labels the initial dyad $\sigma_{\vec{r}}^{(0)} = |L_{r_0}\rangle\langle R_{r_0}|$ sampled in step 3. For $t \geq 1$, the entry r_t gives the index of the Kraus operator chosen at the t th circuit element and we write $\mathcal{K}_{\vec{r}}^{(t)}(\cdot) = K_{r_t}^{(t)}(\cdot)K_{r_t}^{(t)\dagger}$, and use $q_{\vec{r}}^{(t)}$ to denote the corresponding prefactor. Let $\sigma_{\vec{r}}^{(t)}$ denote the current dyad updated up to the t th Kraus operator along the trajectory \vec{r} , so that we have the recursive relation $\sigma_{\vec{r}}^{(t)} = q_{\vec{r}}^{(t)} \mathcal{K}_{\vec{r}}^{(t)}(\sigma_{\vec{r}}^{(t-1)})/P_{\vec{r}}^{(t)}$, where $P_{\vec{r}}^{(t)}$ is the probability of obtaining the outcome corresponding to the map $\mathcal{K}_{\vec{r}}^{(t)}$. The probability $P_{\vec{r}}$ of choosing the trajectory \vec{r} is given by $P_{\vec{r}} = \prod_{t=0}^T P_{\vec{r}}^{(t)}$, where $P_{\vec{r}}^{(0)} = |\alpha_{r_0}|/\|\alpha\|_1$ is the probability of sampling the initial dyad $\sigma_{\vec{r}}^{(0)}$. For $t \geq 1$, $P_{\vec{r}}^{(t)} = \|q_{\vec{r}}^{(t)} \mathcal{K}_{\vec{r}}^{(t)}(\sigma_{\vec{r}}^{(t-1)})\|_1$ is calculated in the t th call to Algorithm 2. Then the final dyad $\sigma_{\vec{r}}^{(T)}$ that we obtain from sampling the trajectory \vec{r} is

$$\sigma_{\vec{r}}^{(T)} = \frac{q_{\vec{r}} \mathcal{K}_{\vec{r}}[\sigma_{\vec{r}}^{(0)}]}{P_{\vec{r}}/P_{\vec{r}}^{(0)}}, \quad (\text{C24})$$

where $\mathcal{K}_{\vec{r}}(\cdot) = \mathcal{K}_{\vec{r}}^{(T)} \circ \dots \circ \mathcal{K}_{\vec{r}}^{(1)}(\cdot)$ and $q_{\vec{r}} = \prod_{t=1}^T q_{\vec{r}}^{(t)}$. This dyad is properly normalized according to the trace norm, but is defined only for those trajectories with nonzero probabilities $P_{\vec{r}}^{(t)} > 0$ for all t . We write \mathcal{R} to denote the set of all such nonzero probability trajectories.

Now, there are two mutually exclusive possibilities for a given iteration of Algorithm 3: either we pick $r_t > 0$ at each circuit element, choose some $\vec{r} \in \mathcal{R}$, and thus obtain a normalized dyad $\sigma_{\vec{r}}^{(T)}$, or at some step we choose $r_t = 0$, and the iteration terminates with $\sigma_{\vec{r}}^{(T)} = 0$. Since these are the only possible outcomes, the total probability of terminating must be $P_{\text{term}} = 1 - \sum_{\vec{r} \in \mathcal{R}} P_{\vec{r}}$. We can now write down an explicit expression for the expectation value of the random variable μ_m in step 15:

$$\langle \mu_m \rangle = P_{\text{term}} \cdot 0 + \sum_{\vec{r} \in \mathcal{R}} P_{\vec{r}} \text{Re}\{\|\alpha\|_1 e^{i\theta_{r_0}} \text{Tr}[\Pi \sigma_{\vec{r}}^{(T)}]\}, \quad (\text{C25})$$

$$= \sum_{\vec{r} \in \mathcal{R}} P_{\vec{r}}^{(0)} \text{Re}\left(\|\alpha\|_1 e^{i\theta_{r_0}} \text{Tr}\{\Pi q_{\vec{r}} \mathcal{K}_{\vec{r}}[\sigma_{\vec{r}}^{(0)}]\}\right), \quad (\text{C26})$$

where in the second line we cancel the factors $P_{\vec{r}}^{(t)}$ for $t \geq 1$ with those in the denominator of Eq. (C24). The real vectors $\vec{r} \notin \mathcal{R}$ are never chosen when running the algorithm, since they correspond to paths where $P_{\vec{r}}^{(t)} = 0$ for some t , and hence $\mathcal{K}_{\vec{r}}^{(t)}[\sigma_{\vec{r}}^{(t-1)}] = 0$. Since $\mathcal{K}_{\vec{r}}[\sigma_{\vec{r}}^{(0)}] = 0$ for all

$\vec{r} \notin \mathcal{R}$, we can add these zero-probability trajectories to the summation, Eq. (C26), without affecting the total. Thus

$$\begin{aligned} \langle \mu_m \rangle &= \sum_{\vec{r}} P_{\vec{r}}^{(0)} \text{Re}\left(\|\alpha\|_1 e^{i\theta_{r_0}} \text{Tr}\{\Pi q_{\vec{r}} \mathcal{K}_{\vec{r}}[\sigma_{\vec{r}}^{(0)}]\}\right), \quad (\text{C27}) \\ &= \sum_{r_0} P_{\vec{r}}^{(0)} \text{Re}\{\|\alpha\|_1 e^{i\theta_{r_0}} \text{Tr} \\ &\quad \times [\Pi \sum_{r_1, \dots, r_T} q_{\vec{r}} \mathcal{K}_{\vec{r}}(|L_{r_0}\rangle\langle R_{r_0}|)]\}, \end{aligned}$$

where in the second line we write $P_{\vec{r}}^{(0)}$ outside of the inner sum since this probability is independent of r_t for $t \geq 1$. The inner expression sums over all Kraus trajectories, and by linearity we have

$$\sum_{r_1, \dots, r_T} q_{\vec{r}} \mathcal{K}_{\vec{r}} = \sum_{r_T} q_{\vec{r}}^{(T)} \mathcal{K}_{\vec{r}}^{(T)} \circ \dots \circ \sum_{r_1} q_{\vec{r}}^{(1)} \mathcal{K}_{\vec{r}}^{(1)}, \quad (\text{C28})$$

$$= O^{(T)} \circ \dots \circ O^{(1)} = \mathcal{E}. \quad (\text{C29})$$

Hence

$$\langle \mu_m \rangle = \sum_{r_0} P_{\vec{r}}^{(0)} \text{Re}\{\|\alpha\|_1 e^{i\theta_{r_0}} \text{Tr}[\Pi \mathcal{E}(|L_{r_0}\rangle\langle R_{r_0}|)]\} \quad (\text{C30})$$

$$= \text{Re}\left\{\text{Tr}\left[\Pi \mathcal{E}\left(\sum_{r_0} \alpha_{r_0} |L_{r_0}\rangle\langle R_{r_0}| \right)\right]\right\} \quad (\text{C31})$$

$$= \text{Tr}[\Pi \mathcal{E}(\rho)], \quad (\text{C32})$$

where in the second line we use the definition $P_{\vec{r}}^{(0)} e^{i\theta_{r_0}} = \alpha_{r_0}/\|\alpha\|_1$. Hence we prove that $\langle \hat{\mu} \rangle = \text{Tr}[\Pi \mathcal{E}(\rho)]$, so $\hat{\mu}$ is an unbiased estimator, with each individual sample satisfying $|\mu_m| \leq \|\alpha\|_1$. We argue above this implies we need $2\|\alpha\|_1^2 \epsilon^{-2} \log(2p_{\text{fail}}^{-1})$ samples [Eq. (C14)]. To generate each sample, we need to make T calls to STABILIZERUPDATE, and we showed in part (i) that each call is computed in $\text{poly}(n)$ time. Therefore, the total runtime is $\|\alpha\|_1^2 \epsilon^{-2} \log(p_{\text{fail}}^{-1}) T \text{poly}(n)$, as stated in Theorem 14.

APPENDIX D: TRACE-NORM ERROR FOR BBCCGH SPARSIFICATION

As discussed in Sec. VII B1, the BBCCGH sparsification lemma [25, Lemma 6] entails that, given a pure state with exact stabilizer decomposition $|\psi\rangle = \sum_j c_j |\phi_j\rangle$, one can randomly generate a k -term sparsification $|\Omega\rangle$, such that

$$\mathbb{E}(\|\psi\rangle - |\Omega\rangle\|^2) \leq \frac{\|c\|_1^2}{k}, \quad (\text{D1})$$

where $\|\cdot\|$ is the standard vector norm. In order to compare with our new sparsification result, which deals with density

operators, we need to translate this in terms of the trace norm. Here we prove the following simple corollary to the BBCCGH sparsification lemma.

Corollary 1. *Given a normalized state $|\psi\rangle = \sum_j c_j |\phi_j\rangle$, for any $k > 0$, one can sample from a distribution of sparsified vectors $|\Omega\rangle = (\|c\|_1/k) \sum_{\alpha=1}^k |\omega_\alpha\rangle$, where $|\omega_\alpha\rangle$ are stabilizer states, such that*

$$\mathbb{E}(\| |\psi\rangle\langle\psi| - |\Omega\rangle\langle\Omega| \|_1) \leq 2 \frac{\|c\|_1}{\sqrt{k}} + \frac{\|c\|_1^2}{k}. \quad (\text{D2})$$

Proof. Let $|\Delta\rangle = |\psi\rangle - |\Omega\rangle$. Then for any particular $|\Omega\rangle$ we have

$$|\psi\rangle\langle\psi| - |\Omega\rangle\langle\Omega| = |\psi\rangle\langle\psi| - (|\psi\rangle\langle\psi| + |\Delta\rangle\langle\Delta| - |\Delta\rangle\langle\psi| - |\psi\rangle\langle\Delta|), \quad (\text{D3})$$

$$= |\Delta\rangle\langle\psi| + |\psi\rangle\langle\Delta| - |\Delta\rangle\langle\Delta|. \quad (\text{D4})$$

Using the triangle inequality

$$\| |\psi\rangle\langle\psi| - |\Omega\rangle\langle\Omega| \|_1 \leq 2 \| |\Delta\rangle\langle\psi| \|_1 + \| |\Delta\rangle\langle\Delta| \|_1, \quad (\text{D5})$$

$$= 2 \| |\Delta\rangle \| \cdot \| |\psi\rangle \| + \| |\Delta\rangle \|^2, \quad (\text{D6})$$

$$= 2 \| |\Delta\rangle \| + \| |\Delta\rangle \|^2, \quad (\text{D7})$$

where the last line follows because $|\psi\rangle$ is normalized. Since the above is true for any $|\Omega\rangle$ taken from the distribution, it follows that

$$\mathbb{E}(\| |\psi\rangle\langle\psi| - |\Omega\rangle\langle\Omega| \|_1) \leq 2\mathbb{E}(\| |\Delta\rangle \|) + \mathbb{E}(\| |\Delta\rangle \|^2). \quad (\text{D8})$$

For the second term, the BBCCGH sparsification lemma [25, Lemma 6] tells us that we have $\mathbb{E}(\| |\Delta\rangle \|^2) \leq \|c\|_1^2/k$.

This leaves the first term. From Jensen's inequality, for any random variable X , we have that $\mathbb{E}(X) \leq \sqrt{\mathbb{E}(X^2)}$. So,

$$\mathbb{E}(\| |\Delta\rangle \|) \leq \sqrt{\mathbb{E}(\| |\Delta\rangle \|^2)}, \quad (\text{D9})$$

$$\leq \frac{\|c\|_1}{\sqrt{k}}, \quad (\text{D10})$$

where the second line again follows from Ref. [25, Lemma 6]. Substituting into the inequality, Eq. (D8), we obtain the result. \blacksquare

APPENDIX E: POSTSELECTION, STABILIZER FIDELITY, AND THE SPARSIFICATION TAIL BOUND IN BBCCGH

In this section we discuss technical difficulties that arise when applying the sparsification results of Bravyi

et al. [25] in the context of the bit-string sampling algorithm. Recall that the main BBCCGH sparsification lemma [25, Lemma 6] only tells us that randomly chosen k -term sparsifications $|\Omega\rangle$ will be close to the target state $|\psi\rangle = \sum_j c_j |\phi_j\rangle$ on average, specifically $\mathbb{E}(\| |\psi\rangle - |\Omega\rangle \|^2) \leq \|c\|_1^2/k$. In itself, this does not preclude the possibility of occasionally obtaining $|\Omega\rangle$ that are very poor estimates for $|\psi\rangle$. One can check numerically that this is a rare occurrence, but it is preferable to put rigorous bounds on the probability of obtaining such outliers. Bravyi *et al.* addressed this with the sparsification tail bound [25, Lemma 7]. This states that if we set $k \geq \|c\|_1^2/\delta^2$, the probability of obtaining $|\Omega\rangle$ close to $|\psi\rangle$ is lower bounded as follows.

$$\Pr \left\{ \| |\psi\rangle - |\Omega\rangle \|^2 \leq \langle\Omega|\Omega\rangle - 1 + \delta^2 \right\} \geq 1 - 2 \exp \left[-\frac{\delta^2}{8F(\psi)} \right], \quad (\text{E1})$$

where $F(\psi)$ is the stabilizer fidelity, defined $F(\psi) = \max_\phi |\langle\phi|\psi\rangle|^2$ where ϕ are stabilizer states. However, there are two subtleties involved in applying this result in practice.

First, note that the usefulness of the bound depends on the norm of $|\Omega\rangle$ being close to (or smaller than) 1. But in general $\langle\Omega|\Omega\rangle$ can be larger. In principle, it is possible for it to be as large as $\|c\|_1^2$, though this is rather unlikely. In any case, this can be solved by using a postselection step where we estimate $\langle\Omega|\Omega\rangle$ (e.g., using FASTNORM) and then discard if we find $\langle\Omega|\Omega\rangle - 1 \gg \delta^2$. Note that normalizing $|\Omega\rangle$ does not solve this problem, as the BBCCGH sparsification results do not tell us about the closeness of $|\psi\rangle$ with $|\Omega\rangle/\| |\Omega\rangle \|$, only the unnormalized vector $|\Omega\rangle$.

Assuming we successfully obtain $|\Omega\rangle$ with sufficiently small norm, a second difficulty arises from the right-hand side of Eq. (E1). The probability of success is larger when the stabilizer fidelity is small. In Bravyi *et al.* [25] it is argued that the failure probability is negligible for cases of interest where stabilizer fidelity is exponentially small in the number of qubits n . Let us unpack this argument by considering a specific case. Assume for the sake of argument that $\langle\Omega|\Omega\rangle$ is very close to 1, so that the expression in square brackets in Eq. (E1) is $\| |\psi\rangle - |\Omega\rangle \|^2 \lesssim \delta^2$. Now suppose we fix target precision δ , and we want to achieve success probability at least p . By rearranging Eq. (E1) we see that this is possible only if the stabilizer fidelity satisfies

$$F(\psi) \leq \frac{\delta^2}{8 \log \left(\frac{2}{1-p} \right)}. \quad (\text{E2})$$

To make this concrete, let us use the modest assumptions that we want trace-norm error δ no larger than

10%, and success probability better than 1/2. This can be achieved only when $F(\psi) \lesssim 0.0009$. Now consider the case where $|\psi\rangle = |T\rangle^{\otimes t}$, where $|T\rangle = (|0\rangle + e^{i\pi/4}|1\rangle)/\sqrt{2}$. For t -fold tensor products of m -qubit states where $m \leq 3$, stabilizer fidelity is multiplicative [25], so that we have $F(|T\rangle^{\otimes t}) \approx (0.854)^t$. It follows that Eq. (E2) is satisfied for these parameters only when we have at least 45 copies of $|T\rangle$. If we want improved accuracy and success probability, the minimum value of t needs to satisfy Eq. (E2) increases. Furthermore, the sparsification tail bound has the curious property that it seems to suggest worse performance for states containing less magic, as quantified by the stabilizer fidelity. For example, if instead of the $\pi/8$ state $|T\rangle$ we consider t -fold tensor products of the $\pi/32$ state $|\pi/32\rangle = (|0\rangle + e^{i\pi/16}|1\rangle)/\sqrt{2}$, we must have at least $t \approx 1200$ before Eq. (E1) gives a nontrivial lower bound on success probability. Therefore, there is a large class of interesting intermediate-sized quantum circuits for which the BBCCGH sparsification tail bound cannot be applied.

Our improved sparsification results in Sec. VII B 2 sidestep these difficulties by considering the difference in the trace norm between $|\psi\rangle\langle\psi|$ and the ensemble $\rho_1 = \sum_{\Omega} \Pr(\Omega) \frac{|\Omega\rangle\langle\Omega|}{\langle\Omega|\Omega\rangle}$ from which sparsified vectors are drawn, rather than fidelity with any particular $|\Omega\rangle$. This allows us to implement classical bit-string sampling with a distribution δ -close to the quantum distribution, even though any particular $|\Omega\rangle$ may not be a good approximation to $|\psi\rangle$. The key idea is that the measurement statistics on the ensemble ρ_1 mimic those on $|\psi\rangle\langle\psi|$; the comparison with any individual $|\Omega\rangle$ is unimportant in the context of bit-string sampling.

APPENDIX F: PROOF OF ENSEMBLE SAMPLING LEMMA

Here we prove Lemma 17, the first of the two lemmata leading to our sparsification result. Given target state $|\psi\rangle = \sum_j c_j |\phi_j\rangle$, we need to prove that, for randomly generated sparse vectors $|\Omega\rangle = (\|c\|_1/k) \sum_{\alpha} |\omega_{\alpha}\rangle$ output from SPARSIFY (Fig. 5), where $|\omega_{\alpha}\rangle$ are stabilizer states randomly drawn from $\{(c_j/|c_j|)|\phi_j\rangle\}$, the following holds:

$$\delta_S = \|\rho_1 - |\psi\rangle\langle\psi|\|_1 \leq \frac{2\|c\|_1^2}{k} + \sqrt{\text{Var}[\langle\Omega|\Omega\rangle]}. \quad (\text{F1})$$

Here k is the number of terms in the sparsified vector $|\Omega\rangle$ and ρ_1 is the ensemble over all possible normalized $|\Omega\rangle$:

$$\rho_1 := \mathbb{E} \left[\frac{|\Omega\rangle\langle\Omega|}{\langle\Omega|\Omega\rangle} \right] = \sum_{\Omega} \Pr(\Omega) \frac{|\Omega\rangle\langle\Omega|}{\langle\Omega|\Omega\rangle}. \quad (\text{F2})$$

First we introduce the operator

$$\rho_2 = \frac{1}{\mu} \mathbb{E} [|\Omega\rangle\langle\Omega|], \quad (\text{F3})$$

where $\mu = \mathbb{E}[\langle\Omega|\Omega\rangle]$. Then using the triangle inequality,

$$\delta_S = \|\rho_1 + \rho_2 - \rho_2 - |\psi\rangle\langle\psi|\|_1, \quad (\text{F4})$$

$$\leq \|\rho_1 - \rho_2\|_1 + \|\rho_2 - |\psi\rangle\langle\psi|\|_1. \quad (\text{F5})$$

Now,

$$\|\rho_1 - \rho_2\|_1 = \left\| \mathbb{E} \left[\frac{|\Omega\rangle\langle\Omega|}{\langle\Omega|\Omega\rangle} \right] - \frac{\mathbb{E} [|\Omega\rangle\langle\Omega|]}{\mu} \right\|_1, \quad (\text{F6})$$

$$= \left\| \mathbb{E} \left[|\Omega\rangle\langle\Omega| \left(\frac{1}{\langle\Omega|\Omega\rangle} - \frac{1}{\mu} \right) \right] \right\|_1. \quad (\text{F7})$$

Using Jensen's inequality we can bring the expectation value outside the norm so that

$$\begin{aligned} \|\rho_1 - \rho_2\|_1 &\leq \mathbb{E} \left\| |\Omega\rangle\langle\Omega| \left(\frac{1}{\langle\Omega|\Omega\rangle} - \frac{1}{\mu} \right) \right\|_1, \\ &= \mathbb{E} \left| \langle\Omega|\Omega\rangle \left(\frac{1}{\langle\Omega|\Omega\rangle} - \frac{1}{\mu} \right) \right|, \\ &= \frac{1}{\mu} \mathbb{E} |\mu - \langle\Omega|\Omega\rangle|. \end{aligned} \quad (\text{F8})$$

That $\mu = \mathbb{E}[\langle\Omega|\Omega\rangle] = 1 + (\|c\|_1^2 - 1)/k$ comes from Ref. [25]. Loosening Eq. (F8) with $\mu^{-1} \leq 1$ gives

$$\|\rho_1 - \rho_2\|_1 \leq \mathbb{E} |\mu - \langle\Omega|\Omega\rangle|, \quad (\text{F9})$$

which is simply the average deviation of $\langle\Omega|\Omega\rangle$ from the mean. Using Jensen's inequality we get

$$\begin{aligned} \mathbb{E} |\mu - \langle\Omega|\Omega\rangle| &\leq \sqrt{\mathbb{E} |\mu - \langle\Omega|\Omega\rangle|^2}, \\ &= \sqrt{\text{Var}[\langle\Omega|\Omega\rangle]}, \end{aligned} \quad (\text{F10})$$

and so

$$\|\rho_1 - \rho_2\|_1 \leq \sqrt{\text{Var}[\langle\Omega|\Omega\rangle]}. \quad (\text{F11})$$

Next, we consider the term $\|\rho_2 - |\psi\rangle\langle\psi|\|_1$, by first finding an explicit form for ρ_2 . Observe that

$$|\Omega\rangle\langle\Omega| = \frac{\|c\|_1^2}{k^2} \sum_{\alpha,\beta} |\omega_{\alpha}\rangle\langle\omega_{\beta}|, \quad (\text{F12})$$

recalling that $|\omega_{\alpha}\rangle = c_j |\phi_j\rangle/|c_j|$ with probability $|c_j|/\|c\|_1$, so that $\mathbb{E}(|\omega_{\alpha}\rangle) = |\psi\rangle/\|c\|_1$ [25]. Taking the expectation

value we have

$$\mathbb{E}[\langle \Omega | \Omega \rangle] = \mu \rho_2 = \frac{\|c\|_1^2}{k^2} \sum_{\alpha, \beta} \mathbb{E}[|\omega_\alpha\rangle\langle\omega_\beta|]. \quad (\text{F13})$$

Let $\sigma := \mathbb{E}[|\omega_\alpha\rangle\langle\omega_\alpha|]$. We split Eq. (F13) into two summations as follows:

$$\mu \rho_2 = \frac{\|c\|_1^2}{k^2} \left[\left(\sum_{\alpha \neq \beta} \mathbb{E}[|\omega_\alpha\rangle\langle\omega_\beta|] \right) + \left(\sum_{\alpha} \mathbb{E}[|\omega_\alpha\rangle\langle\omega_\alpha|] \right) \right], \quad (\text{F14})$$

$$= \frac{\|c\|_1^2}{k^2} \left[\left(\sum_{\alpha \neq \beta} \frac{|\psi\rangle\langle\psi|}{\|c\|_1^2} \right) + \left(\sum_{\alpha} \sigma \right) \right], \quad (\text{F15})$$

$$= \frac{1}{k^2} \left(\sum_{\alpha \neq \beta} |\psi\rangle\langle\psi| \right) + \frac{\|c\|_1^2}{k^2} \sum_{\alpha} \sigma.$$

In the first contribution, we use the independence of ω_α and ω_β when $\alpha \neq \beta$ and $\mathbb{E}[|\omega_\alpha\rangle] = |\psi\rangle/\|c\|_1$. Next, there are $k(k-1)$ terms and k terms in the first and second summations, respectively, so that

$$\mu \rho_2 = (1 - k^{-1}) |\psi\rangle\langle\psi| + \frac{\|c\|_1^2}{k} \sigma. \quad (\text{F16})$$

Using this form for ρ_2 , we have that

$$\begin{aligned} \|\rho_2 - |\psi\rangle\langle\psi|\|_1 &= \mu^{-1} \|\mu \rho_2 - \mu |\psi\rangle\langle\psi|\|_1, \quad (\text{F17}) \\ &= \mu^{-1} \|(1 - k^{-1} - \mu) |\psi\rangle\langle\psi| \\ &\quad + \frac{\|c\|_1^2}{k} \sigma\|_1. \end{aligned}$$

Substituting in the value of μ we find $1 - k^{-1} - \mu = -\|c\|_1^2/k$ and so

$$\|\rho_2 - |\psi\rangle\langle\psi|\|_1 = \frac{\|c\|_1^2}{k\mu} \|\sigma - |\psi\rangle\langle\psi|\|_1 \leq 2 \frac{\|c\|_1^2}{k}, \quad (\text{F18})$$

where we use the triangle inequality, $\|\sigma\|_1 = 1$ and $\mu^{-1} \leq 1$. Substituting Eq. (F11) and Eq. (F18) into Eq. (F5), completes the proof of the lemma.

APPENDIX G: SPARSIFICATION VARIANCE BOUND

We now prove Lemma 18, the second lemma leading to Theorem 16. Recall that given a state $|\psi\rangle = \sum_j c_j |\phi_j\rangle$,

where $|\phi_j\rangle$ are stabilizer states, we can obtain a sparsified k -term approximation given by

$$|\Omega\rangle = \frac{\|c\|_1}{k} \sum_{\alpha=1}^k |\omega_\alpha\rangle, \quad (\text{G1})$$

where each $|\omega_\alpha\rangle$ is chosen randomly so that $|\omega_\alpha\rangle = (c_j/|c_j|) |\phi_j\rangle$ with probability $p_j = |c_j|/\|c\|_1$. In general $|\Omega\rangle$ may not be conventionally normalized, but Lemma 18 upper bounds the variance of $\langle \Omega | \Omega \rangle$. We now prove Lemma 18.

Proof of Lemma 18. In Ref. [25] it was shown that

$$\mu = \mathbb{E}[\langle \Omega | \Omega \rangle] = \frac{\|c\|_1^2}{k} + \frac{\|c\|_1^2}{k^2} \mathbb{E}(B), \quad (\text{G2})$$

where $B = \sum_{\alpha} \sum_{\beta \neq \alpha} \langle \omega_\alpha | \omega_\beta \rangle$. Since $|\omega_\alpha\rangle$ and $|\omega_\beta\rangle$ are independently sampled for distinct α and β , we get

$$\mathbb{E}(\langle \omega_\alpha | \omega_\beta \rangle) = \mathbb{E}(\langle \omega_\alpha |) \mathbb{E}(|\omega_\beta\rangle) = \frac{\langle \psi | \psi \rangle}{\|c\|_1^2}. \quad (\text{G3})$$

We use similar proof techniques to bound $\mathbb{E}[\langle \Omega | \Omega \rangle^2]$, and in turn bound the variance. We begin with

$$\langle \Omega | \Omega^2 \rangle = \frac{\|c\|_1^4}{k^4} \left(\sum_{\alpha, \beta} \langle \omega_\alpha | \omega_\beta \rangle \right)^2, \quad (\text{G4})$$

$$= \frac{\|c\|_1^4}{k^4} \left[\sum_{\alpha} \left(\langle \omega_\alpha | + \sum_{\beta \neq \alpha} \langle \omega_\alpha | \omega_\beta \rangle \right) \right]^2, \quad (\text{G5})$$

$$= \frac{\|c\|_1^4}{k^4} (k + B)^2, \quad (\text{G6})$$

$$= \frac{\|c\|_1^4}{k^4} (k^2 + 2kB + B^2), \quad (\text{G7})$$

where in the second line we note that there are k terms in the summation. Whereas from Eq. (G2) we have

$$\mathbb{E}[\langle \Omega | \Omega \rangle]^2 = \frac{\|c\|_1^4}{k^4} [k^2 + 2k\mathbb{E}(B) + \mathbb{E}(B)^2]. \quad (\text{G8})$$

Comparing these expressions, for the variance we obtain

$$\text{Var}[\langle \Omega | \Omega \rangle] = \mathbb{E}[\langle \Omega | \Omega \rangle^2] - \mathbb{E}[\langle \Omega | \Omega \rangle]^2, \quad (\text{G9})$$

$$= \frac{\|c\|_1^4}{k^4} [\mathbb{E}(B^2) - \mathbb{E}(B)^2]. \quad (\text{G10})$$

By counting terms in the summation B , and using the relation, Eq. (G3), we find

$$\mathbb{E}(B)^2 = \frac{k^2(k-1)^2}{\|c\|_1^4}. \quad (\text{G11})$$

Expanding B^2 , we get

$$B^2 = \left(\sum_{\alpha} \sum_{\beta \neq \alpha} \langle \omega_{\alpha} | \omega_{\beta} \rangle \right) \left(\sum_{\lambda} \sum_{\mu \neq \lambda} \langle \omega_{\lambda} | \omega_{\mu} \rangle \right) \quad (\text{G12})$$

$$= \sum_{(\alpha, \beta, \lambda, \mu) \in \mathcal{A}} \langle \omega_{\alpha} | \omega_{\beta} \rangle \langle \omega_{\lambda} | \omega_{\mu} \rangle + B', \quad (\text{G13})$$

where \mathcal{A} denotes the set of all possible combinations $(\alpha, \beta, \lambda, \mu)$ where all four indices are distinct, and B' denotes the remaining terms where at least two of the indices are the same. Now, if $(\alpha, \beta, \lambda, \mu)$ are all distinct, then $\langle \omega_{\alpha} | \omega_{\beta} \rangle$ and $\langle \omega_{\lambda} | \omega_{\mu} \rangle$ are independent random variables, so $\mathbb{E}(\langle \omega_{\alpha} | \omega_{\beta} \rangle \langle \omega_{\lambda} | \omega_{\mu} \rangle) = \mathbb{E}(\langle \omega_{\alpha} | \omega_{\beta} \rangle) \mathbb{E}(\langle \omega_{\lambda} | \omega_{\mu} \rangle)$. This yields

$$\mathbb{E}(B^2) = \frac{k(k-1)(k-2)(k-3)}{\|c\|_1^4} + \mathbb{E}(B'). \quad (\text{G14})$$

Substituting Eqs. (G11) and (G14) back into Eq. (G10), we obtain

$$\text{Var}[\langle \Omega | \Omega \rangle] = \frac{\|c\|_1^4}{k^4} \mathbb{E}(B') - \frac{k(k-1)(4k-6)}{k^4}. \quad (\text{G15})$$

We now consider terms $\langle \omega_{\alpha} | \omega_{\beta} \rangle \langle \omega_{\lambda} | \omega_{\mu} \rangle$ in the expansion of B^2 where $(\alpha, \beta, \lambda, \mu)$ are not all distinct. We use the notation $B_{j=k}$ to indicate the sum of all terms where indices j and k are equal but all others are distinct, e.g., $B_{\lambda=\alpha} = \sum_{\alpha, \beta, \mu} \langle \omega_{\alpha} | \omega_{\beta} \rangle \langle \omega_{\alpha} | \omega_{\mu} \rangle$, where the summation is over terms such that $\alpha, \beta,$ and μ are all distinct, and so on. There are $k(k-1)(k-2)$ terms in each summation of this type. Similarly for the terms sharing two pairs of indices, we use the notation $B_{\lambda=\alpha; \mu=\beta} = \sum_{\alpha \neq \beta} \langle \omega_{\alpha} | \omega_{\beta} \rangle \langle \omega_{\alpha} | \omega_{\beta} \rangle$. These summations comprise of $k(k-1)$ terms. From Eq. (G12), we never have terms where $\alpha = \beta$ or $\lambda = \mu$. We can therefore write

$$B' = B_{\lambda=\alpha} + B_{\mu=\alpha} + B_{\lambda=\beta} + B_{\mu=\beta} + B_{\lambda=\alpha; \mu=\beta} + B_{\mu=\alpha; \lambda=\beta}. \quad (\text{G16})$$

One can check that $\mathbb{E}[B_{\lambda=\alpha}^*] = \mathbb{E}[B_{\mu=\beta}]$ and $\mathbb{E}[B_{\mu=\alpha}^*] = \mathbb{E}[B_{\lambda=\beta}]$. Therefore,

$$\mathbb{E}[B'] = 2\text{Re}\{\mathbb{E}[B_{\lambda=\beta}] + \mathbb{E}[B_{\mu=\beta}]\} + \mathbb{E}[B_{\lambda=\alpha; \mu=\beta} + B_{\mu=\alpha; \lambda=\beta}]. \quad (\text{G17})$$

Next we note that

$$\mathbb{E}[B_{\lambda=\beta}] = \sum_{\alpha} \sum_{\beta \neq \alpha} \sum_{\alpha \neq \mu \neq \beta} \mathbb{E}[\langle \omega_{\alpha} | \omega_{\beta} \rangle \langle \omega_{\beta} | \omega_{\mu} \rangle], \quad (\text{G18})$$

$$= k(k-1)(k-2) \mathbb{E}[\langle \omega_{\alpha} | \omega_{\beta} \rangle \langle \omega_{\beta} | \omega_{\mu} \rangle], \quad (\text{G19})$$

$$= \frac{k(k-1)(k-2)}{\|c\|_1^2} \langle \psi | \sigma | \psi \rangle,$$

where $\sigma = \mathbb{E}[|\omega_{\beta}\rangle\langle\omega_{\beta}|] = \sum_j (|c_j|/\|c\|_1) |\phi_j\rangle\langle\phi_j|$, since the probability of sampling $|\omega_{\beta}\rangle\langle\omega_{\beta}| = |\phi_j\rangle\langle\phi_j|$ is defined as $p_j = |c_j|/\|c\|_1$. Next we consider $\mathbb{E}[B_{\mu=\beta}]$. Taking the modulus and using the triangle inequality we obtain

$$|\mathbb{E}[B_{\mu=\beta}]| \leq \sum_{\alpha} \sum_{\beta \neq \alpha} \sum_{\alpha \neq \lambda \neq \beta} \mathbb{E}[|\langle \omega_{\alpha} | \omega_{\beta} \rangle \langle \omega_{\lambda} | \omega_{\beta} \rangle|], \quad (\text{G20})$$

$$= k(k-1)(k-2) \mathbb{E}[|\langle \omega_{\alpha} | \omega_{\beta} \rangle \langle \omega_{\beta} | \omega_{\lambda} \rangle|], \quad (\text{G21})$$

$$= \frac{k(k-1)(k-2)}{\|c\|_1^2} \langle \psi | \sigma | \psi \rangle. \quad (\text{G22})$$

Similarly, for the last two terms $B'' = B_{\lambda=\alpha; \mu=\beta} + B_{\mu=\alpha; \lambda=\beta}$, we obtain

$$|\mathbb{E}[B'']| \leq \sum_{\alpha} \sum_{\beta \neq \alpha} \mathbb{E}[|\langle \omega_{\alpha} | \omega_{\beta} \rangle \langle \omega_{\alpha} | \omega_{\beta} \rangle|], \quad (\text{G23})$$

$$+ \sum_{\alpha} \sum_{\beta \neq \alpha} \mathbb{E}[\langle \omega_{\alpha} | \omega_{\beta} \rangle \langle \omega_{\beta} | \omega_{\alpha} \rangle], \quad (\text{G24})$$

$$= 2 \sum_{\alpha} \sum_{\beta \neq \alpha} \mathbb{E}[\langle \omega_{\alpha} | \omega_{\beta} \rangle \langle \omega_{\beta} | \omega_{\alpha} \rangle]. \quad (\text{G25})$$

Using cyclicity of the trace get

$$\mathbb{E}[\langle \omega_{\alpha} | \omega_{\beta} \rangle \langle \omega_{\beta} | \omega_{\alpha} \rangle] = \mathbb{E}[\text{Tr}[\langle \omega_{\alpha} | \omega_{\beta} \rangle \langle \omega_{\beta} | \omega_{\alpha} \rangle]], \quad (\text{G26})$$

$$= \text{Tr}[\mathbb{E}[|\omega_{\alpha}\rangle\langle\omega_{\alpha}|] \mathbb{E}[|\omega_{\beta}\rangle\langle\omega_{\beta}|]], \quad (\text{G27})$$

$$= \text{Tr}[\mathbb{E}[|\omega_{\alpha}\rangle\langle\omega_{\alpha}|] \mathbb{E}[|\omega_{\beta}\rangle\langle\omega_{\beta}|]], \quad (\text{G28})$$

$$= \text{Tr}[\sigma^2], \quad (\text{G29})$$

so that

$$|\mathbb{E}[B'']| \leq 2k(k-1) \text{Tr}[\sigma^2] \leq 2k(k-1). \quad (\text{G30})$$

Combining the results, Eqs. (G17), (G19), (G22), and (G30), gives us

$$\mathbb{E}[B'] \leq 4 \frac{k(k-1)(k-2)}{\|c\|_1^2} \langle \psi | \sigma | \psi \rangle + 2k(k-1). \quad (\text{G31})$$

Writing

$$C = \|c\|_1^2 \langle \psi | \sigma | \psi \rangle = \|c\|_1 \sum_j |c_j| |\langle \psi | \phi_j \rangle|^2 \quad (\text{G32})$$

and substituting the expression for $\mathbb{E}(B')$ into Eq. (G15) we obtain

$$\text{Var}[\langle \Omega | \Omega \rangle] \leq 4 \frac{k^3 - 3k^2 + 2k}{k^4} C + 2 \frac{\|c\|_1^4}{k^2} \left(1 - \frac{1}{k}\right) - \frac{4k^3 - 10k^2 + 6k}{k^4}, \quad (\text{G33})$$

which to leading order in $1/k$ is

$$\text{Var}[\langle \Omega | \Omega \rangle] \leq \frac{4(C-1)}{k} + 2 \left(\frac{\|c\|_1^2}{k} \right)^2 + \mathcal{O}\left(\frac{C}{k^3}\right), \quad (\text{G34})$$

which gives us the general bound appearing in Lemma 18.

Clifford magic states are defined in Ref. [25] as those pure states $|\psi\rangle$ that are stabilized by a group \mathcal{Q} of Clifford unitary operators whose generators take the form $UX_j U^\dagger$, where X_j is the Pauli X operator that acts on the j th qubit. For such states, there exists [25] an optimal decomposition

$$|\psi\rangle = \sum_{q \in \mathcal{Q}} c_q |\phi_q\rangle = \frac{1}{|\mathcal{Q}| \langle \psi | \phi_0 \rangle} \sum_{q \in \mathcal{Q}} q |\phi_0\rangle, \quad (\text{G35})$$

where $|\phi_0\rangle$ is some stabilizer state that achieves the maximize possible value for $|\langle \psi | \phi_0 \rangle|$. If we take this decomposition as the basis for our sparsification, then we have

$$\|c\|_1 = |\mathcal{Q}| \cdot (|\mathcal{Q}| |\langle \psi | \phi_0 \rangle|)^{-1} = |\langle \psi | \phi_0 \rangle|^{-1} \quad (\text{G36})$$

and

$$\sigma = \sum_{q \in \mathcal{Q}} p_q q |\phi_0\rangle \langle \phi_0| q^\dagger, \quad (\text{G37})$$

where $p_q = |\mathcal{Q}|^{-1}$. This yields

$$\langle \psi | \sigma | \psi \rangle = \sum_{q \in \mathcal{Q}} p_q \langle \psi | q |\phi_0\rangle \langle \phi_0| q^\dagger | \psi \rangle, \quad (\text{G38})$$

$$= \sum_{q \in \mathcal{Q}} p_q \langle \psi | \phi_0 \rangle \langle \phi_0 | \psi \rangle, \quad (\text{G39})$$

$$= |\langle \psi | \phi_0 \rangle|^2 = \frac{1}{\|c\|_1^2}, \quad (\text{G40})$$

where in the second line we use the Hermiticity of q and $q|\psi\rangle = |\psi\rangle$. This shows that for optimal decompositions of Clifford magic states, $C = 1$, and leads to the simplified bound

$$\text{Var}[\langle \Omega | \Omega \rangle] \leq 2 \left(\frac{\|c\|_1^2}{k} \right)^2 + \frac{2}{k^3}. \quad (\text{G41})$$

■

Finally, we comment on the effect of the constant C when $|\psi\rangle$ is not a Clifford magic state. Recall that C can be written in terms of this expected overlap, $C = \|c\|_1^2 \mathbb{E}[|\langle \psi | \omega \rangle|^2]$, and enters into Theorem 16 via the critical precision $\delta_c = 8(C-1)/\|c\|_1^2$. Consider $|\psi\rangle = |\psi'\rangle^{\otimes N}$ where $|\psi'\rangle$ are pure states. When $|\psi\rangle$ is a product of N pure states, we can write each




FIG. 8. The variable C as introduced in Eq. (G32) as a function of the angle θ for a class of single-qubit states. This is the C value for one copy of the state, for n copies we must raise to the n th power. The prefactor $C-1$ appears in Eq. (G33) and is important because when $C=1$, the variance scales asymptotically as $\mathcal{O}(1/k^2)$. We highlight two specific angles $\theta = \{\pi/8, 0.1187\}$ that correspond to angles used in Fig. 6. For $\theta = \pi/8$, we have $C-1=0$ and so the $\mathcal{O}(1/k^2)$ is exact as can be seen in Fig. 6. For $\theta = 0.1187$, we have the maximal possible value of C and Fig. 6 shows the maximal deviation from $\mathcal{O}(1/k^2)$ scaling.

randomly sampled stabilizer state as $|\omega\rangle = \bigotimes_{\alpha=1}^N |\omega_\alpha\rangle$, where $|\omega_\alpha\rangle$ are i.i.d. random vectors. It follows that $\mathbb{E}[|\langle \psi | \omega \rangle|^2] = (\mathbb{E}[|\langle \psi' | \omega_\alpha \rangle|^2])^N$. Since $|\omega_\alpha\rangle$ are always stabilizer states, when $|\psi'\rangle$ are nonstabilizer states, we have $|\langle \psi' | \omega_\alpha \rangle|^2 < 1$. Therefore, the threshold precision $\delta_c < 8C/\|c\|_1^2 = 8(\mathbb{E}[|\langle \psi' | \omega_\alpha \rangle|^2])^N$ vanishes for large N when $|\psi\rangle$ is a tensor product of N pure states. Moreover, in Fig. 8 we plot values of C for a class of single-qubit states, showing that $C-1$ is close to zero even when N is not large.

APPENDIX H: BIT-STRING SAMPLING SIMULATOR TECHNICAL DETAILS

In this appendix, we give full pseudocode for our bit-string sampling simulator (Algorithm 4), prove its validity as a method to classically emulate sampling from the quantum distribution $P(\vec{x}) = \text{Tr}[\Pi_{\vec{x}} \rho]$, and analyze its runtime. This constitutes a proof of Theorem 19. As described in the main text, Algorithm 4 draws bit strings \vec{x} from a classical distribution $P_{\text{sim}}(\vec{x})$, using two subroutines from Ref. [25], SPARSIFY and FASTNORM. As sketched in the main text, our strategy is to define an idealized algorithm EXACT where calls to FASTNORM are replaced by an oracle, which can compute $\|\Pi_{\vec{y}} |\Omega\rangle\|$ exactly for any unnormalized $|\Omega\rangle$ and bit string \vec{y} . The algorithm EXACT draws from a distribution $P_{\text{ex}}(\vec{x})$. We first show that P_{ex} is δ_S -close to the quantum distribution P . We then argue that the distribution P_{sim} that Algorithm 4 draws from is ϵ -close to P_{ex} . Finally we optimize the choice of δ_S and ϵ and analyze the runtime.

Input: Decomposition $\rho = \sum_j p_j |\psi_j\rangle\langle\psi_j|$, where for each ψ_j we have a known stabilizer state decomposition $|\psi_j\rangle = \sum_r c_r^{(j)} |\phi_r^{(j)}\rangle$. Real numbers δ_S, p_{FN} and ϵ_{FN} . Number of bits w .

Output: String \vec{x} of length w , sampled from a distribution $P''(\vec{x}) = \sum_j p_j P_j(\vec{x})$, which approximates $P(\vec{x}) = \text{Tr}(\Pi_{\vec{x}}\rho)$, where $\Pi_{\vec{x}} = |\vec{x}\rangle\langle\vec{x}| \otimes \mathbb{1}_{n-w}$.

- 1: Select index j with probability p_j .
- 2: $k \leftarrow \lceil 4\|c^{(j)}\|_1^2 \delta_S^{-1} \rceil$
- 3: $|\Omega'\rangle \leftarrow \text{SPARSIFY}(|\psi_j\rangle, k)$
- 4: $W \leftarrow \text{FASTNORM}(|\Omega'\rangle, p_{\text{FN}}, \epsilon_{\text{FN}})$
- 5: $|\Omega\rangle \leftarrow |\Omega'\rangle/\sqrt{W}$
- 6: $\vec{x} \leftarrow ()$ (initialize empty string)
- 7: $P_{\vec{x}} \leftarrow 1$
- 8: **for** $b \leftarrow 1$ **to** w **do**
- 9: $P_{(\vec{x},0)} \leftarrow \text{FASTNORM}(\Pi_{(\vec{x},0)}|\Omega\rangle, p_{\text{FN}}, \epsilon_{\text{FN}})$
- 10: $P(x_b = 0|\vec{x}) \leftarrow P_{(\vec{x},0)}/P_{\vec{x}}$
- 11: **if** $P_{x_b=0} < 1/2$ **then**
- 12: $P(x_b = 1|\vec{x}) \leftarrow 1 - P_{x_b=0}$
- 13: **else**
- 14: $P_{(\vec{x},1)} \leftarrow \text{FASTNORM}(\Pi_{(\vec{x},1)}|\Omega\rangle, p_{\text{FN}}, \epsilon_{\text{FN}})$
- 15: $P(x_b = 1|\vec{x}) \leftarrow P_{(\vec{x},1)}/P_{\vec{x}}$
- 16: $P(x_b = 0|\vec{x}) \leftarrow 1 - P_{x_b=1}$
- 17: **end if**
- 18: Select $y \in \{0, 1\}$ with probability $P(x_b = y|\vec{x})$, then $x_b \leftarrow y$.
- 19: $P_{(\vec{x},x_b)} \leftarrow P_{\vec{x}} \times P(x_b|\vec{x})$
- 20: $\vec{x} \leftarrow (\vec{x}, x_b)$
- 21: **end for**
- 22: **return** \vec{x}

Algorithm 4. Bit-string sampling for mixed states.

EXACT is identical to our Algorithm 4, except where our algorithm estimates probabilities $\|\Pi_{\vec{y}}|\Omega\rangle\|^2$ using FASTNORM, EXACT computes them exactly. Therefore EXACT first samples a state $|\psi_j\rangle$ from the ensemble with probability p_j , and chooses a sparsification $|\Omega\rangle = |\Omega_{j,l}\rangle$ with probability $q_{j,l} = \text{Pr}(\Omega_{j,l}|\psi_j)$. Given the selected $|\Omega\rangle$, a bit string is sampled by choosing each bit in turn via a series of conditional probabilities:

$$\text{Pr}(\vec{x}|\Omega) = \text{Pr}(x_1)\text{Pr}(x_2|\vec{x}_1) \dots \text{Pr}(x_w|\vec{x}_{w-1}), \quad (\text{H1})$$

$$= \frac{\|\Pi_{\vec{x}}|\Omega\rangle\|^2}{\|\Omega\|^2} = \text{Tr} \left[\Pi_{\vec{x}} \frac{|\Omega\rangle\langle\Omega|}{\langle\Omega|\Omega\rangle} \right]. \quad (\text{H2})$$

Here we use the notation \vec{x}_m to denote the string comprised of the first m bits of \vec{x} , so that $\Pi_{\vec{x}_m} = \bigotimes_{j=1}^m |x_j\rangle\langle x_j| \otimes \mathbb{1}_{n-m}$. We take \vec{x}_0 to be the empty string, so that $\Pi_{\vec{x}_0} = \mathbb{1}$. The probability of choosing $y \in \{0, 1\}$ for the m th bit, given $m-1$ bits already sampled, is computed as

$$\text{Pr}(y|\vec{x}_{m-1}) = \frac{\|\Pi_{(x_1, \dots, x_{m-1}, y)}|\Omega\rangle\|^2}{\|\Pi_{\vec{x}_{m-1}}|\Omega\rangle\|}. \quad (\text{H3})$$

Thus EXACT outputs bit strings \vec{x} sampled from a distribution:

$$P_{\text{ex}}(\vec{x}) = \sum_j \sum_l p_j q_{j,l} \frac{\|\Pi_{\vec{x}}|\Omega_{j,l}\rangle\|^2}{\|\Omega_{j,l}\|^2}, \quad (\text{H4})$$

$$\begin{aligned} &= \sum_j \sum_l p_j q_{j,l} \frac{\text{Tr}[\Pi_{\vec{x}}|\Omega_{j,l}\rangle\langle\Omega_{j,l}|]}{\langle\Omega_{j,l}|\Omega_{j,l}\rangle}, \\ &= \text{Tr} \left[\Pi_{\vec{x}} \sum_j p_j \sum_l q_{j,l} \frac{|\Omega_{j,l}\rangle\langle\Omega_{j,l}|}{\langle\Omega_{j,l}|\Omega_{j,l}\rangle} \right], \\ &= \text{Tr} \left[\Pi_{\vec{x}} \sum_j p_j \mathbb{E} \left(\frac{|\Omega\rangle\langle\Omega|}{\langle\Omega|\Omega\rangle} \right) \right] = \text{Tr}[\Pi_{\vec{x}}\rho'], \end{aligned}$$

where $\rho' = \sum_j p_j \rho_1^{(j)}$, and each $\rho_1^{(j)}$ given by

$$\rho_1^{(j)} := \sum_{\Omega} \text{Pr}(\Omega|\psi_j) \frac{|\Omega\rangle\langle\Omega|}{\langle\Omega|\Omega\rangle}. \quad (\text{H5})$$

In other words $\rho_1^{(j)}$ is the expected sparsification given target pure state $|\psi_j\rangle$, as defined in Eq. (103). In step 2, k is chosen so that by Theorem 16, we have $\|\rho_1^{(j)} - |\psi_j\rangle\langle\psi_j|\|_1 \leq \delta_S + \mathcal{O}(\delta_S^2)$, provided $\delta_S \geq \delta_c$, where δ_c is the critical precision. We return to the $\delta_S < \delta_c$ case at the end of this appendix. By the triangle inequality we have

$$\|\rho' - \rho\|_1 = \left\| \sum_j p_j \rho_1^{(j)} - \sum_j p_j |\psi_j\rangle\langle\psi_j| \right\|_1, \quad (\text{H6})$$

$$\leq \sum_j p_j \|\rho_1^{(j)} - |\psi_j\rangle\langle\psi_j|\|_1, \quad (\text{H7})$$

$$\leq \sum_j p_j [\delta_S + \mathcal{O}(\delta_S^2)] = \delta_S + \mathcal{O}(\delta_S^2). \quad (\text{H8})$$

Since $P_{\text{ex}}(\vec{x}) = \text{Tr}[\Pi_{\vec{x}}\rho']$ and for the quantum distribution we have $P(\vec{x}) = \text{Tr}[\Pi_{\vec{x}}\rho]$, It follows that $\|P_{\text{ex}} - P\|_1 \leq \delta_S + \mathcal{O}(\delta_S^2)$.

It remains to show that using a sequence of calls to FASTNORM, Algorithm 4 generates probability distributions $P_{\text{sim}}(\vec{x})$ that well approximate $P_{\text{ex}}(\vec{x})$, where

$$P_{\text{sim}}(\vec{x}) = \sum_j p_j q_{j,l} Q_{j,l}(\vec{x}). \quad (\text{H9})$$

Here each $Q_{j,l}(\vec{x})$ is the probability of Algorithm 4 returning \vec{x} given the sparsification $|\Omega_{j,l}\rangle$. We now drop the subscript as we consider a single sparsification $|\Omega\rangle$. Recall that FASTNORM takes as input error parameters p_{FN} and ϵ_{FN} , and unnormalized vectors $\Pi_{\vec{y}}|\Omega\rangle$ with known k -term stabilizer decomposition. Then with probability $(1 - p_{\text{FN}})$

it outputs a random variable η that approximates $\|\Pi_{\vec{y}}|\Omega\rangle\|^2$ to within a multiplicative error of ϵ_{FN} :

$$(1 - \epsilon_{\text{FN}})\|\Pi_{\vec{y}}|\Omega\rangle\|^2 \leq \eta \leq (1 + \epsilon_{\text{FN}})\|\Pi_{\vec{y}}|\Omega\rangle\|^2. \quad (\text{H10})$$

Algorithm 4 approximates the chain of conditional probabilities H1 by calls to FASTNORM. The probability of choosing $y \in \{0, 1\}$ for the m th bit of \vec{x} , conditioned on the first $m - 1$ bits being \vec{x}_{m-1} is therefore bounded as

$$\epsilon_- \frac{\|\Pi_{(\vec{x}_{m-1}, y)}|\Omega\rangle\|^2}{\|\Pi_{\vec{x}_{m-1}}|\Omega\rangle\|^2} \leq \Pr(y|\vec{x}_{m-1}) \leq \epsilon_+ \frac{\|\Pi_{(\vec{x}_{m-1}, y)}|\Omega\rangle\|^2}{\|\Pi_{\vec{x}_{m-1}}|\Omega\rangle\|^2},$$

with probability $(1 - p_{\text{FN}})^2$, where

$$\epsilon_{\pm} = \frac{1 \pm \epsilon_{\text{FN}}}{1 \mp \epsilon_{\text{FN}}}. \quad (\text{H11})$$

So, given a particular sparsification $|\Omega\rangle$, the w bit string \vec{x} is sampled from a distribution $\mathcal{Q}(\vec{x})$, which satisfies

$$\prod_{m=1}^w \epsilon_- \frac{\|\Pi_{\vec{x}_m}|\Omega\rangle\|^2}{\|\Pi_{\vec{x}_{m-1}}|\Omega\rangle\|^2} \leq \mathcal{Q}(\vec{x}) \leq \prod_{m=1}^w \epsilon_+ \frac{\|\Pi_{\vec{x}_m}|\Omega\rangle\|^2}{\|\Pi_{\vec{x}_{m-1}}|\Omega\rangle\|^2}$$

with probability at least $(1 - p_{\text{FN}})^{2w}$. This simplifies to

$$\frac{(1 - \epsilon_{\text{FN}})^w \|\Pi_{\vec{x}}|\Omega\rangle\|^2}{(1 + \epsilon_{\text{FN}})^w \|\Omega\rangle\|^2} \leq \mathcal{Q}(\vec{x}) \leq \frac{(1 + \epsilon_{\text{FN}})^w \|\Pi_{\vec{x}}|\Omega\rangle\|^2}{(1 - \epsilon_{\text{FN}})^w \|\Omega\rangle\|^2}. \quad (\text{H12})$$

One can check that $(1 + \epsilon_{\text{FN}})^w / (1 - \epsilon_{\text{FN}})^w \leq 1 + 3w\epsilon_{\text{FN}}$, whenever $\epsilon_{\text{FN}} \leq 1/5$, and the analogous result holds for the lower bound. Therefore, $\mathcal{Q}_{j,l}(\vec{x})$ approximates $\|\Pi_{\vec{x}}|\Omega_{j,l}\rangle\|^2 / \|\Omega_{j,l}\rangle\|^2$ up to multiplicative error $3w\epsilon_{\text{FN}}$. Comparing Eq. (H4) with Eq. (H9), we therefore obtain

$$(1 - 3w\epsilon_{\text{FN}})P_{\text{ex}}(\vec{x}) \leq P_{\text{sim}}(\vec{x}) \leq (1 + 3w\epsilon_{\text{FN}})P_{\text{ex}}(\vec{x}). \quad (\text{H13})$$

If we want to bound the total multiplicative error due to the sequence of calls to FASTNORM to ϵ , then we must set $\epsilon_{\text{FN}} = \epsilon / (3w)$. It then follows that

$$\|P_{\text{sim}} - P_{\text{ex}}\|_1 \leq \epsilon. \quad (\text{H14})$$

In the first part of the proof we show that $\|P_{\text{ex}} - P\|_1 \leq \delta_S + \mathcal{O}(\delta_S^2)$ (provided we are above the critical precision threshold δ_c). Combined with Eq. (H14), we obtain

$$\|P_{\text{sim}} - P\|_1 \leq \epsilon + \delta_S + \mathcal{O}(\delta_S^2), \quad (\text{H15})$$

where $P(\vec{x}) = \text{Tr}[\Pi_{\vec{x}}\rho]$.

Similarly the error bound given above is only obtained with probability $(1 - p_{\text{FN}})^{2w} \approx 1 - 2wp_{\text{FN}}$, so to obtain

the above closeness in ℓ_1 norm, with failure probability at most p_{fail} , we must set $p_{\text{FN}} = p_{\text{fail}} / (2w)$. If we select the state $|\psi_j\rangle$ in step 1, then $k \leq 4\|c^{(j)}\|_1^2 \delta_S^{-1} + 1$. To return a single bit string \vec{x} there are at most $2w$ calls to FASTNORM, so the runtime is $\mathcal{O}(wkn^3 \epsilon_{\text{FN}}^{-2} \log p_{\text{FN}}^{-1}) = \mathcal{O}[w^3 n^3 \|c^{(j)}\|_1^2 \delta_S^{-1} \epsilon^2 \log(w/p_{\text{fail}})]$. Recall that the statement of the theorem defined the quantity $\tilde{\Xi} = \sum_j p_j \|c^{(j)}\|_1^2$, so that the time T to obtain a single bit string is non-deterministic. The *expected* (average-case) runtime is $\mathcal{O}[w^3 n^3 \tilde{\Xi} \delta_S^{-1} \epsilon^2 \log(w/p_{\text{fail}})]$. If the decomposition is optimal with respect to the monotone Ξ , then we have $\tilde{\Xi} = \Xi(\rho)$ and the average-case runtime is $\mathcal{O}[\Xi(\rho)]$. For equimagical states, $\Xi(\rho) = \xi(\psi_j)$ for all j , and this expression becomes the worst-case runtime.

We now optimize the choice of δ_S and ϵ . Setting the total error budget $\delta = \delta_S + \epsilon$, by inspecting the runtime we find that the best constant is obtained by setting $\delta_S = \delta/3$ and $\epsilon = 2\delta/3$. The constraint $\delta_S \geq 8D$ therefore becomes $\delta \geq 24D$. Substituting the optimal choice of δ_S and ϵ into the expected runtime, we obtain

$$\mathbb{E}(T) = \mathcal{O}[w^3 n^3 \tilde{\Xi} \delta^{-3} \log(w/p_{\text{fail}})]. \quad (\text{H16})$$

The above holds for the case where the sparsification error δ_S is no smaller than a critical value $\delta_c = 8(C_j - 1) / \|c^{(j)}\|_1^2$, where $C_j = \|c^{(j)}\|_1 \sum_r |c_r| |\langle \psi_j | \phi_r \rangle|^2$ is defined for the randomly chosen pure state $|\psi_j\rangle$. Therefore, to ensure we are above the critical error regime for any $|\psi_j\rangle$, we can require that $\delta_S \geq 8D$, where $D = \max\{(C_j - 1) / \|c^{(j)}\|_1^2\}$. This entails $\delta \geq 24D$ for the overall precision.

Now suppose that we want to achieve arbitrary precision, $\delta < 24D$. In this regime, one can amend the expression for k in step 2 to achieve any desired precision, at the cost of slightly poorer scaling in the runtime. We first use Lemmas 17 and 18 to obtain a sharpened bound on the sparsification error:

$$\delta_S \leq 2 \frac{\|c^{(j)}\|_1^2}{k} + \sqrt{\frac{\|c^{(j)}\|_1^2}{k} \sqrt{4D + 2 \frac{\|c^{(j)}\|_1^2}{k} + \mathcal{O}\left(\frac{1}{k^2}\right)}}. \quad (\text{H17})$$

When $\delta_S \ll 8D$, we can achieve a precision of δ_S by choosing

$$k \approx 4\|c^{(j)}\|_1^2 \left(\frac{D}{\delta_S^2} + \frac{1}{\delta_S} \right) + \mathcal{O}(1). \quad (\text{H18})$$

Substituting the revised expression for k into the expected runtime, with $\delta_S = \delta/3$ and $\epsilon = 2\delta/3$, we obtain

$$\mathbb{E}(T) = \mathcal{O}[w^3 n^3 \tilde{\Xi} (\delta^{-3} + 3D\delta^{-4}) \log(w/p_{\text{fail}})]. \quad (\text{H19})$$

Here we recover the same asymptotic δ^{-3} scaling as derived from the original BBCCGH sparsification

lemma [25]. However, the prefactor from this prior work was two, whereas our prefactor D is typically exponentially small in the number of qubits (see Appendix G). Therefore, at intermediate precision, the δ^{-4} term may still dominate. When the target precision δ is too small, our bound on the required k exceeds the number of terms in the exact decomposition of $|\psi\rangle$ [i.e., the decomposition achieving the stabilizer rank $\chi(\psi)$]. In this scenario, using a sparsified approximation in both our approach and in Ref. [25] has no benefit, and one should instead use an exact decomposition without any sparsification.

-
- [1] R. P. Feynman, Simulating physics with computers, *Int. J. Theor. Phys.* **21**, 467 (1982).
- [2] D. Gottesman, Theory of fault-tolerant quantum computation, *Phys. Rev. A* **57**, 127 (1998).
- [3] S. Aaronson and D. Gottesman, Improved simulation of stabilizer circuits, *Phys. Rev. A* **70**, 052328 (2004).
- [4] L. G. Valiant, Quantum circuits that can be simulated classically in polynomial time, *SIAM J. Comput.* **31**, 1229 (2002).
- [5] B. M. Terhal and D. P. DiVincenzo, Classical simulation of noninteracting-fermion quantum circuits, *Phys. Rev. A* **65**, 032325 (2002).
- [6] R. Jozsa and A. Miyake, Matchgates and classical simulation of quantum circuits, *Proc. R. Soc. Lond. A* **464**, 3089 (2008).
- [7] D. J. Brod, Efficient classical simulation of matchgate circuits with generalized inputs and measurements, *Phys. Rev. A* **93**, 062332 (2016).
- [8] Y.-Y. Shi, L.-M. Duan, and G. Vidal, Classical simulation of quantum many-body systems with a tree tensor network, *Phys. Rev. A* **74**, 022320 (2006).
- [9] M. Schwarz and M. V. d. Nest, Simulating quantum circuits with sparse output distributions, [arXiv:1310.6749](https://arxiv.org/abs/1310.6749) (2013).
- [10] M. Oszmaniec and D. J. Brod, Classical simulation of photonic linear optics with lost particles, *New J. Phys.* **20**, 092002 (2018).
- [11] M. J. Bremner, A. Montanaro, and D. J. Shepherd, Achieving quantum supremacy with sparse and noisy commuting quantum computations, *Quantum* **1**, 8 (2017).
- [12] M. Van Den Nest, Simulating quantum computers with probabilistic methods, *Quantum Info. Comput.* **11**, 784 (2011).
- [13] H. De Raedt, F. Jin, D. Willsch, M. Willsch, N. Yoshioka, N. Ito, S. Yuan, and K. Michielsen, Massively parallel quantum computer simulator, eleven years later, *Comput. Phys. Commun.* **237**, 47 (2019).
- [14] F. Arute *et al.*, Quantum supremacy using a programmable superconducting processor, *Nature* **574**, 505 (2019).
- [15] R. S. Bennink, E. M. Ferragut, T. S. Humble, J. A. Laska, J. J. Nutaro, M. G. Pleszkoch, and R. C. Pooser, Unbiased simulation of near-Clifford quantum circuits, *Phys. Rev. A* **95**, 062337 (2017).
- [16] K. Temme, S. Bravyi, and J. M. Gambetta, Error Mitigation for Short-Depth Quantum Circuits, *Phys. Rev. Lett.* **119**, 180509 (2017).
- [17] S. Bravyi and A. Kitaev, Universal quantum computation with ideal Clifford gates and noisy ancillas, *Phys. Rev. A* **71**, 022316 (2005).
- [18] E. T. Campbell, B. M. Terhal, and C. Vuillot, Roads towards fault-tolerant universal quantum computation, *Nature* **549**, 172 (2017).
- [19] D. Stahlke, Quantum interference as a resource for quantum speedup, *Phys. Rev. A* **90**, 022302 (2014).
- [20] H. Pashayan, J. J. Wallman, and S. D. Bartlett, Estimating Outcome Probabilities of Quantum Circuits Using Quasiprobabilities, *Phys. Rev. Lett.* **115**, 070501 (2015).
- [21] M. Howard and E. Campbell, Application of a Resource Theory for Magic States to Fault-Tolerant Quantum Computing, *Phys. Rev. Lett.* **118**, 090501 (2017).
- [22] J. R. Seddon and E. T. Campbell, Quantifying magic for multi-qubit operations, *Proc. R. Soc. A* **475**, 20190251 (2019).
- [23] S. Bravyi, G. Smith, and J. A. Smolin, Trading Classical and Quantum Computational Resources, *Phys. Rev. X* **6**, 021043 (2016).
- [24] S. Bravyi and D. Gosset, Improved Classical Simulation of Quantum Circuits Dominated by Clifford Gates, *Phys. Rev. Lett.* **116**, 250501 (2016).
- [25] 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).
- [26] H. Qassim, J. J. Wallman, and J. Emerson, Clifford recompilation for faster classical simulation of quantum circuits, *Quantum* **3**, 170 (2019).
- [27] C. Ferrie and J. Emerson, Frame representations of quantum mechanics and the necessity of negativity in quasiprobability representations, *J. Phys. A: Math. Theor.* **41**, 352001 (2008).
- [28] P. Rall, D. Liang, J. Cook, and W. Kretschmer, Simulation of qubit quantum circuits via Pauli propagation, *Phys. Rev. A* **99**, 062337 (2019).
- [29] U. Leonhardt, Discrete Wigner function and quantum-state tomography, *Phys. Rev. A* **53**, 2998 (1996).
- [30] D. Gross, Hudson's theorem for finite-dimensional quantum systems, *J. Math. Phys.* **47**, 122107 (2006).
- [31] V. Veitch, C. Ferrie, D. Gross, and J. Emerson, Negative quasi-probability as a resource for quantum computation, *New J. Phys.* **14**, 113011 (2012).
- [32] E. T. Campbell, H. Anwar, and D. E. Browne, Magic-State Distillation in all Prime Dimensions Using Quantum Reed-Muller Codes, *Phys. Rev. X* **2**, 041021 (2012).
- [33] H. Anwar, B. J. Brown, E. T. Campbell, and D. E. Browne, Fast decoders for qudit topological codes, *New J. Phys.* **16**, 063038 (2014).
- [34] E. T. Campbell, Enhanced Fault-Tolerant Quantum Computing in D-level Systems, *Phys. Rev. Lett.* **113**, 230501 (2014).
- [35] V. Veitch, S. H. Mousavian, D. Gottesman, and J. Emerson, The resource theory of stabilizer quantum computation, *New J. Phys.* **16**, 013009 (2014).
- [36] Or equivalently additive $[\mathcal{M}(\rho \otimes \sigma) = \mathcal{M}(\rho) + \mathcal{M}(\sigma)]$ after taking the logarithm.
- [37] S. Mansfield and E. Kashefi, Quantum Advantage from Sequential-Transformation Contextuality, *Phys. Rev. Lett.* **121**, 230401 (2018).

- [38] N. Delfosse, P. A. Guerin, J. Bian, and R. Raussendorf, Wigner Function Negativity and Contextuality in Quantum Computation on Rebits, *Phys. Rev. X* **5**, 021003 (2015).
- [39] R. Raussendorf, D. E. Browne, N. Delfosse, C. Okay, and J. Bermejo-Vega, Contextuality and Wigner-function negativity in qubit quantum computation, *Phys. Rev. A* **95**, 052334 (2017).
- [40] R. Raussendorf, J. Bermejo-Vega, E. Tyhurst, C. Okay, and M. Zurek, Phase-space-simulation method for quantum computation with magic states on qubits, *Phys. Rev. A* **101**, 012350 (2020).
- [41] M. Heinrich and D. Gross, Robustness of magic and symmetries of the stabiliser polytope, *Quantum* **3**, 132 (2019).
- [42] L. Kocia and P. Love, Stationary phase method in discrete Wigner functions and classical simulation of quantum circuits, [arXiv:1810.03622](https://arxiv.org/abs/1810.03622) (2018).
- [43] Y. Huang and P. Love, Approximate stabilizer rank and improved weak simulation of Clifford-dominated circuits for qudits, *Phys. Rev. A* **99**, 052307 (2019).
- [44] L. Kocia and M. Sarovar, Improved simulation of quantum circuits by fewer gaussian eliminations, *Phys. Rev. A* **103**, 022603 (2020).
- [45] H. Pashayan, S. D. Bartlett, and D. Gross, From estimation of quantum probabilities to simulation of quantum circuits, *Quantum* **4**, 223 (2020).
- [46] A. Heimendahl, F. Montealegre-Mora, F. Vallentin, and D. Gross, Stabilizer extent is not multiplicative, [arXiv:2007.04363](https://arxiv.org/abs/2007.04363) (2020).
- [47] K. Fang and Z.-W. Liu, No-Go Theorems for Quantum Resource Distillation, *Phys. Rev. Lett.* **125**, 060405 (2020).
- [48] S. D. Bartlett, B. C. Sanders, S. L. Braunstein, and K. Nemoto, Efficient Classical Simulation of Continuous Variable Quantum Information Processes, *Phys. Rev. Lett.* **88**, 097904 (2002).
- [49] A. Mari and J. Eisert, Positive Wigner Functions Render Classical Simulation of Quantum Computation Efficient, *Phys. Rev. Lett.* **109**, 230503 (2012).
- [50] E. Chitambar and G. Gour, Quantum resource theories, *Rev. Mod. Phys.* **91**, 025001 (2019).
- [51] T. Baumgratz, M. Cramer, and M. B. Plenio, Quantifying Coherence, *Phys. Rev. Lett.* **113**, 140401 (2014).
- [52] A. Streltsov, G. Adesso, and M. B. Plenio, Quantum coherence as a resource, *Rev. Mod. Phys.* **89**, 041003 (2017).
- [53] P. O. Boykin, T. Mor, M. Pulver, V. Roychowdhury, and F. Vatan, On universal and fault-tolerant quantum computing, [arXiv:quant-ph/9906054](https://arxiv.org/abs/quant-ph/9906054) [quant-ph] (1999).
- [54] O. Rudolph, A new class of entanglement measures, *J. Math. Phys.* **42**, 5306 (2001).
- [55] B. Regula, Convex geometry of quantum resource quantification, *J. Phys. A: Math. Theor.* **51**, 045303 (2017).
- [56] 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).
- [57] G. Vidal and R. Tarrach, Robustness of entanglement, *Phys. Rev. A* **59**, 141 (1999).
- [58] O. Rudolph, Further results on the cross norm criterion for separability, *Quant. Inf. Proc.* **4**, 219 (2005).
- [59] M. Steiner, Generalized robustness of entanglement, *Phys. Rev. A* **67**, 054305 (2003).
- [60] S. Boyd and L. Vandenberghe, *Convex Optimization* (Cambridge University Press, Cambridge 2004).
- [61] A. Uhlmann, Roofs and convexity, *Entropy* **12**, 1799 (2010).
- [62] K. G. H. Vollbrecht and R. F. Werner, Entanglement measures under symmetry, *Phys. Rev. A* **64**, 062307 (2001).
- [63] M. B. Plenio, Logarithmic Negativity: A Full Entanglement Monotone That Is Not Convex, *Phys. Rev. Lett.* **95**, 090503 (2005).
- [64] X. Wang, M. M. Wilde, and Y. Su, Efficiently Computable Bounds for Magic State Distillation, *Phys. Rev. Lett.* **124**, 090505 (2020).
- [65] E. T. Campbell, Catalysis and activation of magic states in fault-tolerant architectures, *Phys. Rev. A* **83**, 032317 (2011).
- [66] 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).
- [67] M. Beverland, E. Campbell, M. Howard, and V. Kliuchnikov, Lower bounds on the non-Clifford resources for quantum computations, [arXiv:1904.01124](https://arxiv.org/abs/1904.01124) (2019).
- [68] F. G. S. L. Brandão and G. Gour, Reversible Framework for Quantum Resource Theories, *Phys. Rev. Lett.* **115**, 070503 (2015).
- [69] N. Datta, Min- and max-relative entropies and a new entanglement monotone, *IEEE Trans. Inf. Theory* **55**, 2816 (2009).
- [70] W. Hoeffding, Probability inequalities for sums of bounded random variables, *J. Am. Stat. Assoc.* **58**, 13 (1963).
- [71] J. J. Wallman and J. Emerson, Noise tailoring for scalable quantum computation via randomized compiling, *Phys. Rev. A* **94**, 052325 (2016).
- [72] E. Campbell, Shorter gate sequences for quantum computing by mixing unitaries, *Phys. Rev. A* **95**, 042306 (2017).
- [73] M. B. Hastings, Turning gate synthesis errors into incoherent errors, *Quantum Info. Comput.* **17**, 488 (2017).
- [74] E. Campbell, Random Compiler for Fast Hamiltonian Simulation, *Phys. Rev. Lett.* **123**, 070503 (2019).
- [75] R. Takagi, B. Regula, K. Bu, Z.-W. Liu, and G. Adesso, Operational Advantage of Quantum Resources in Sub-channel Discrimination, *Phys. Rev. Lett.* **122**, 140402 (2019).
- [76] R. Takagi and B. Regula, General Resource Theories in Quantum Mechanics and Beyond: Operational Characterization via Discrimination Tasks, *Phys. Rev. X* **9**, 031053 (2019).
- [77] Z.-W. Liu, K. Bu, and R. Takagi, One-Shot Operational Quantum Resource Theory, *Phys. Rev. Lett.* **123**, 020401 (2019).
- [78] S. Rana, P. Parashar, A. Winter, and M. Lewenstein, Logarithmic coherence: Operational interpretation of ℓ_1 -norm coherence, *Phys. Rev. A* **96**, 052336 (2017).
- [79] E. Chitambar and G. Gour, Critical Examination of Incoherent Operations and a Physically Consistent Resource Theory of Quantum Coherence, *Phys. Rev. Lett.* **117**, 030401 (2016).

- [80] A. Winter and D. Yang, Operational Resource Theory of Coherence, *Phys. Rev. Lett.* **116**, 120404 (2016).
- [81] K. Ben Dana, M. García Díaz, M. Mejatty, and A. Winter, Resource theory of coherence: Beyond states, *Phys. Rev. A* **95**, 062327 (2017).
- [82] B. Yadin, J. Ma, D. Girolami, M. Gu, and V. Vedral, Quantum Processes Which Do Not Use Coherence, *Phys. Rev. X* **6**, 041028 (2016).
- [83] C. Napoli, T. R. Bromley, M. Cianciaruso, M. Piani, N. Johnston, and G. Adesso, Robustness of Coherence: An Operational and Observable Measure of Quantum Coherence, *Phys. Rev. Lett.* **116**, 150502 (2016).
- [84] C. H. Bennett, D. P. DiVincenzo, C. A. Fuchs, T. Mor, E. Rains, P. W. Shor, J. A. Smolin, and W. K. Wootters, Quantum nonlocality without entanglement, *Phys. Rev. A* **59**, 1070 (1999).
- [85] A. Heimendahl, M. Heinrich, and D. Gross, The axiomatic and the operational approaches to resource theories of magic do not coincide, [arXiv:2011.11651](https://arxiv.org/abs/2011.11651)