General Entropic Constraints on Calderbank-Shor-Steane Codes within Magic Distillation Protocols

Rhea Alexander,[1,2] Si Gvirtz-Chen,[2,*] Nikolaos Koukoulekidis,[1] and David Jennings[1,2]

[1] *Department of Physics, Imperial College London, London SW7 2AZ, United Kingdom*
[2] *School of Physics and Astronomy, University of Leeds, Leeds LS2 9JT, United Kingdom*

(Received 5 January 2023; accepted 30 May 2023; published 30 June 2023)

Magic states are fundamental building blocks on the road to fault-tolerant quantum computing. Calderbank-Shor-Steane (CSS) codes play a crucial role in the construction of magic distillation protocols. Previous work has cast quantum computing with magic states for odd dimension d within a phase-space setting in which universal quantum computing is described by the statistical mechanics of quasiprobability distributions. Here we extend this framework to the important $d = 2$ qubit case and show that we can exploit common structures in CSS circuits to obtain distillation bounds capable of outperforming previous monotone bounds in regimes of practical interest. Moreover, in the case of CSS-code projections, we arrive at a novel cutoff result on the code length n of the CSS code in terms of parameters characterizing a desired distillation, which implies that for fixed target error rate and acceptance probability, one needs to consider only CSS codes below a threshold number of qubits. These entropic constraints are not due simply to the data-processing inequality but rely explicitly on the stochastic representation of such protocols.

DOI: 10.1103/PRXQuantum.4.020359

I. INTRODUCTION

Work towards achieving fault-tolerant quantum computing is currently seeing rapid progress on qubit systems on many computational platforms [1–9]. In particular, the surface code [10–12] is a leading framework that allows Clifford operations to be implemented transversally on blocks of physical qubits. However, Clifford operations are not universal for quantum computing [13,14], and in fact it is impossible to encode any universal gateset transversally [15]. A prominent method of circumventing this problem is the magic state injection model, wherein the Clifford group is promoted to a universal gateset by injecting copies of special nonstabilizer states known as magic states [16,17]. While magic states can only be prepared in surface codes with relatively high error rates, it is possible to reduce the noise per copy by converting many noisy magic states into fewer higher-fidelity magic states using only stabilizer operations [18]. This process, known as *magic state distillation*, allows magic states to be produced with arbitrarily high purity, and thereby enables universal quantum computation within surface-code models.

Almost all protocols [18–22] to date for qubit magic distillation are based on a subclass of stabilizer codes known as Calderbank-Shor-Steane (CSS) codes [23,24]. CSS codes can be constructed from two classical linear codes, allowing one to draw on a plethora of results from classical coding theory to construct quantum codes with desirable properties. For instance, it has been shown that CSS codes are optimal when it comes to constructing quantum error-correcting codes that support a transversal T gate [25], a key feature in many of the aforementioned distillation protocols. Although significant progress has been made to reduce the overhead of such protocols [26], distillation is still estimated to dominate the total resource cost of performing a computation in the magic state injection model. Therefore, a better understanding of the extent to which this cost can be reduced is of great practical interest.

Recent work [27] has developed a framework for analyzing magic distillation in odd-dimensional systems by taking key insights from a rich literature of majorization theory and applying them to discrete phase-space representations of magic states. In odd dimensions, Gross's Wigner function [28] provides a representation wherein distillable magic states correspond to quasiprobability distributions containing negativity on a discrete phase space [16]. By contrast, stabilizer states correspond to probability distributions, and stabilizer operations in general are represented by stochastic transformations [28,29]. Thus when computation is restricted to the stabilizer setting, one obtains a classical stochastic model that can be studied

[*] pysc@leeds.ac.uk

using entropic theory and, in particular, relative majorization [30–34]. Reference [27] extended majorization tools to negatively represented magic states, and found that a dense subset of α-Rényi entropies H_α remain well defined and meaningful as quantifiers of disorder on quasiprobability distributions under stochastic processing, leading to fundamental constraints on magic distillation protocols in the form of thermodynamic laws.

Since most quantum algorithms are formulated for systems of qubits, the important question of whether this framework can be extended to qubit systems remains. There are, however, many well-known obstacles in constructing valid Wigner representations for qubits (related to the fact that 2^{-1} does not exist modulo 2 [29]). Many constructions for Wigner functions, including Gross's, cannot be extended to qubits [35,36], while others represent some pure stabilizer states negatively [37,38], which breaks the link established in odd dimensions between Wigner negativity and quantum computational speedup [39,40]. While substantial work has been done to develop phase-space representations wherein all qubit stabilizer operations are non-negatively represented [41,42], channels are typically not mapped to linear, let alone stochastic, transformations under such representations.

Progress can be made by identifying subsets of qubit stabilizer operations that can be represented stochastically, while nevertheless remaining capable of universal quantum computation via magic state injection. In this paper, we make use of a Wigner representation for qubits introduced in Ref. [43] that shares many desirable features with Gross's representation, such as the linear representation of channels. Drawing on results from Ref. [44], we show that CSS circuits—the subset of stabilizer circuits wherein CSS states play the role of stabilizer states—remain stochastic in this representation. Since CSS circuits are known to be capable of universal quantum computation via magic state injection [44,45], they provide a setting where we can extend the statistical mechanics framework for universal quantum computing developed in Ref. [27] to qubits.

The structure of our paper is as follows. In Sec. III, we introduce the phase-space representation of qubit states and channels we use, and identify some regimes where this representation becomes stochastic. Building from this, in Sec. IV we develop majorization techniques to analyze stochastically represented magic distillation protocols on qubits. Finally, in Sec. V, we apply these tools to derive general entropic constraints (in the form of upper and lower bounds on code length) for distillation protocols that project onto CSS codes, which exploit structures basic to this distillation strategy.

II. MAIN RESULTS

We show that CSS circuits can be represented by stochastic maps on a well-defined multiqubit phase space. By exploiting techniques from majorization, in Theorem 3 we extend the statistical mechanical framework of Ref. [27] to CSS qubit quantum computation.

We further find that, similar to Ref. [47], every CSS circuit can be decomposed in terms of protocols that project onto CSS codes, which therefore constitute the core machinery for magic distillation in CSS circuits. For such CSS-code-projection protocols, we obtain novel upper bounds on the code length n as a function of the number of output qubits k, acceptance probability p and input and output error rates ϵ and δ, respectively [which are typically related to the code distance D via $\delta = O(\epsilon^D)$]. Our main result is the following, which we generalize to odd-dimensional systems and arbitrary stabilizer codes in Theorem 5.

Result 1.—Consider the distillation of k copies of a pure qubit magic state ψ from a supply of the noisy magic state ρ, where both ψ and ρ have real density matrices in the computational basis. Any magic distillation protocol that projects onto the codespace of an $[[n,k]]$ CSS code and can use n copies of ρ to distil out a k-qubit state ρ' at output error $\delta \geq \|\rho' - \psi^{\otimes k}\|_1$ and acceptance probability p must have a code length n such that

$$n \geq \frac{k\left[1 - H_\alpha(W_\psi)\right] - \frac{\alpha}{1-\alpha}\log\left(\frac{p}{1+\delta 2^{5/2}}\right)}{\left[1 - H_\alpha(W_\rho)\right]} \tag{1}$$

for all $\alpha \in \mathcal{A}$ for which $H_\alpha(W_\rho) < 1$, and

$$n \leq \frac{k\left[H_\alpha(W_\psi) - 1\right] + \frac{\alpha}{1-\alpha}\log\left(\frac{p}{1+\delta 2^{5/2}}\right)}{\left[H_\alpha(W_\rho) - 1\right]}, \tag{2}$$

for all $\alpha \in \mathcal{A}$ for which $H_\alpha(W_\rho) > 1$, where $H_\alpha(W_\rho)$ is a Rényi entropy measure computed on the Wigner representation of the quantum state ρ defined in Eq. (8).

Combining these bounds with prior work on projective robustness of magic [46], we constrain the code length of any $[[n,k]]$ CSS-code projection that could achieve a desired distillation process to lie within a finite range. An example for the distillation of a single Hadamard state are shown in Fig. 1. In this case, our analytic upper bounds on the number of noisy magic states needed to distil out a single copy at the output take on the particularly simple form:

$$n \leq \log_{f(\epsilon)}\left[\frac{1+6\delta}{p}\right]^2 =: n^*, \tag{3}$$

where the base of the logarithm is given by $f(\epsilon) := [1 - \epsilon + \epsilon^2/2]^{-1} \geq 1$. In the spirit of Ref. [48], Eq. (3) and more generally Theorem 5, can be viewed as expressing trade-off relations between various distillation parameters. These fundamental no-go results may be instructive when constructing stabilizer-code-projection protocols with optimized parameters.

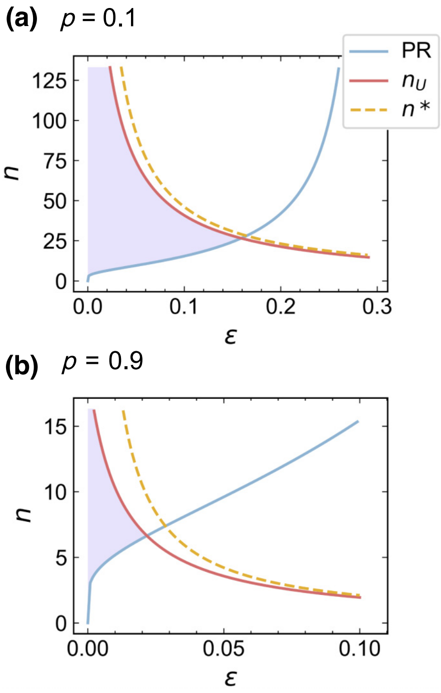

FIG. 1. (Finite range on CSS-code lengths for magic state distillation protocols.) We plot upper and lower bounds on the number of copies n of the noisy Hadamard state $(1-\epsilon)\,|H\rangle\,\langle H| + \epsilon\frac{\mathbb{1}}{2}$ required to distil a single output qubit $|H\rangle$ with output error rate $\delta = 10^{-9}$ by projecting onto an $[[n,1]]$ CSS code. The shaded purple region shows the range of code lengths allowed by the tightest numeric upper bound (red curve) from Theorem 4 and the lower bound from projective robustness (PR) introduced in Ref. [46] (blue curve). The analytic upper bound n^* (dashed yellow curve) defined in Eq. (2) is shown to form a good approximation to the numeric bound. (a) When target acceptance probability p is low ($p = 0.1$) the upper bounds are less constraining. (b) By increasing to $p = 0.9$, the upper bounds become considerably tighter. In both cases, there is a cutoff input error ϵ beyond which no CSS-code-projection protocol can achieve the desired combination of output error and acceptance probability.

III. STOCHASTIC REPRESENTATION FOR CSS CIRCUITS ON QUBITS

In this section, we review the qubit representation W_ρ introduced in Ref. [43] that forms the backbone of our work. We expand upon its properties and confirm that it respects all sequential and parallel composition of processes, the former of which crucially gives rise to a well-defined input-output relation $W_{\mathcal{E}(\rho)} = W_{\mathcal{E}} W_\rho$. Moreover, the latter implies that the representation of product states factorizes over subsystems, a property which will prove computationally advantageous given that inputs to magic distillation protocols typically take on the form $\rho^{\otimes n}$. Furthermore, we show that all magic distillation protocols executed by CSS circuits are stochastically represented, which means that such protocols can be analyzed using majorization theory and admit a description in terms of classical statistical mechanics on quasiprobability distributions.

A. Phase-space representation of qubit states

We first establish some convenient notation. Let $\mathbf{u} := (u_1, \ldots, u_n) \in \mathbb{Z}_2^n$ denote a binary vector. Furthermore, given any single qubit operator O, let us denote

$$O(\mathbf{u}) := O^{u_1} \otimes \cdots \otimes O^{u_n}. \tag{4}$$

With this in place, consider an n-qubit quantum system with total Hilbert space $\mathcal{H}_2^n := \mathcal{H}_2^{\otimes n}$. We associate to this system a phase space $\mathcal{P}_n := \mathbb{Z}_2^n \times \mathbb{Z}_2^n$, where $\mathcal{P}_n$ consists of all vectors $(\mathbf{u}_x, \mathbf{u}_z)$, and has a symplectic inner product $[\mathbf{u}, \mathbf{v}]$ defined as

$$[\mathbf{u}, \mathbf{v}] := \mathbf{u}_z \cdot \mathbf{v}_x - \mathbf{v}_z \cdot \mathbf{u}_x \equiv \mathbf{u}_z \cdot \mathbf{v}_x + \mathbf{v}_z \cdot \mathbf{u}_x, \tag{5}$$

where arithmetic is carried out modulo 2.

We are now in a position to define our chosen representation over $\mathcal{P}_n$ (we refer the reader to Appendix A for further details and proofs of the properties presented). We first define n-qubit displacement operators $\{D_\mathbf{u}\}$, where $\mathbf{u} := (\mathbf{u}_x, \mathbf{u}_z) \in \mathcal{P}_n$, via strings of single qubit Pauli operators X and Z as

$$D_\mathbf{u} := Z(\mathbf{u}_z) X(\mathbf{u}_x), \tag{6}$$

which generate the Heisenberg-Weyl group $H(2)^{\times n}$ on n-qubits modulo phase factors [49]. These displacement operators satisfy

$$D_\mathbf{u} D_\mathbf{v} = (-1)^{[\mathbf{u},\mathbf{v}]} D_\mathbf{v} D_\mathbf{u}. \tag{7}$$

Using these displacement operators, we can construct the following representation:

$$W_\rho(\mathbf{u}) := \frac{1}{2^n} \mathrm{tr}[A_\mathbf{u}^\dagger \rho] \tag{8}$$

for any n-qubit state ρ, where $\{A_\mathbf{u}\}$ are the set of 2^{2n} *phase-point operators* on n qubits, which are defined as

$$A_\mathbf{u} = \frac{1}{2^n} \sum_{\mathbf{v} \in \mathcal{P}_n} (-1)^{[\mathbf{u},\mathbf{v}]} D_\mathbf{v}. \tag{9}$$

It can be shown (see Appendix A 1) that these phase-point operators share the following properties with those defining Gross's Wigner representation of n qudits with Hilbert-space dimension d on a phase space $\mathcal{P} := \mathbb{Z}_d^n \times \mathbb{Z}_d^n$:

(A1) $A_{\mathbf{u}_X \oplus \mathbf{u}_Y} = A_{\mathbf{u}_X} \otimes A_{\mathbf{u}_Y}$ on a bipartite system XY, where $\mathbf{u}_X$ and $\mathbf{u}_Y$ are, respectively, points in the phase spaces of subsystems X and Y.

(A2) $\mathrm{tr}[A_{\mathbf{u}}^{\dagger}A_{\mathbf{v}}] = d^n\delta_{\mathbf{u},\mathbf{v}}$,
(A3) $\mathrm{tr}[A_{\mathbf{u}}] = 1$,
(A4) $\sum_{\mathbf{u}\in\mathcal{P}} A_{\mathbf{u}} = d^n\mathbb{1}^{\otimes n}$.

These properties imply that W_ρ provides an informationally complete and normalized representation of general n-qubit states, i.e.,

$$\sum_{\mathbf{u}} W_\rho(\mathbf{u}) = 1, \tag{10}$$

for any quantum state ρ. Like Gross's Wigner function, an immediate consequence of Eq. (7) is that the representation W_ρ transforms covariantly under the displacement operators, namely,

$$W_{D_{\mathbf{v}}^{\dagger}\rho D_{\mathbf{v}}}(\mathbf{u}) = W_\rho(\mathbf{u}+\mathbf{v}), \tag{11}$$

for all $\mathbf{u},\mathbf{v}\in\mathcal{P}_n$. In fact, everything in the construction of this representation has proceeded in direct analogy to Gross's, except for the lack of phase factors ensuring the Hermiticity of the displacement operators. As a result, the phase-point operators in Eq. (9) are no longer Hermitian, which in turn implies that W_ρ is generally complex.

However, it turns out that the real and imaginary parts of W_ρ are related to the quantum state in the following simple way (a proof is given in Appendix A 2):

Lemma 1.—Given any n-qubit quantum state ρ,

$$\mathrm{Re}[W_\rho(\mathbf{u})] = W_{\mathrm{Re}(\rho)}(\mathbf{u}) \tag{12}$$

$$\mathrm{Im}[W_\rho(\mathbf{u})] = W_{\mathrm{Im}(\rho)}(\mathbf{u}) \tag{13}$$

for all $\mathbf{u}\in\mathcal{P}_n$, where $\mathrm{Re}(\rho)$ and $\mathrm{Im}(\rho)$ are, respectively, the real and imaginary parts of the density matrix of ρ in the computational basis.

This immediately implies the following.

Corollary 1.—The representation W_ρ of an n-qubit state ρ is a quasiprobability distribution if and only if ρ is an n-rebit state, i.e., the density matrix ρ is real in the computational basis.

To simplify our analysis, we, therefore, focus on rebit states for the majority of this work, although in Sec. VI we show how we can handle arbitrary qubit states by treating the real and imaginary components separately resulting in an overcomplete quasiprobability representation. Typically, however, we consider the case of distilling the following Hadamard state

$$|H\rangle := \cos\frac{\pi}{8}|0\rangle + \sin\frac{\pi}{8}|1\rangle, \tag{14}$$

which is equivalent to the canonical magic state $|A\rangle := T|+\rangle = 1/\sqrt{2}(|0\rangle + e^{i(\pi)/4}|1\rangle$ up to a Clifford unitary [50], where $T := \mathrm{diag}(1, e^{i(\pi)/4})$ is the T gate. The Hadamard state can thus be used in a stabilizer gadgetization circuit to implement the T gate [18].

B. Phase-space representation of channels

The representation of qubit states induces a corresponding representation of qubit channels. Let $\mathcal{E}$ be an arbitrary channel from n to m qubits, and

$$\mathcal{J}(\mathcal{E}) = (\mathcal{I}\otimes\mathcal{E})(|\phi_n^+\rangle\langle\phi_n^+|) \tag{15}$$

be its associated Choi state [51], where $|\phi_n^+\rangle$ is the canonical maximally entangled state on two copies of the input system. We now define a representation [52] of a quantum channel $\mathcal{E}$ as

$$W_{\mathcal{E}}(\mathbf{v}|\mathbf{u}) := 2^{2n}W_{\mathcal{J}(\mathcal{E})}(\mathbf{u}\oplus\mathbf{v}), \tag{16}$$

for all $\mathbf{v}\in\mathcal{P}_m$, and $\mathbf{u}\in\mathcal{P}_n$. Under this representation, every channel becomes a matrix mapping the representation of an input state to the representation of the output state. More precisely, if $\sigma = \mathcal{E}(\rho)$, then

$$W_\sigma(\mathbf{v}) = \sum_{\mathbf{u}\in\mathcal{P}_n} W_{\mathcal{E}}(\mathbf{v}|\mathbf{u})W_\rho(\mathbf{u}), \tag{17}$$

for all $\mathbf{v}\in\mathcal{P}_m$.

Furthermore, the representation $W_{\mathcal{E}}$ respects the sequential and parallel composition of channels, i.e.,

$$W_{\mathcal{E}\circ\mathcal{F}} = W_{\mathcal{E}}W_{\mathcal{F}}, \tag{18}$$

$$W_{\mathcal{E}\otimes\mathcal{F}} = W_{\mathcal{E}}\otimes W_{\mathcal{F}}. \tag{19}$$

One useful implication of Eq. (19) is that when $\mathcal{E}$ and $\mathcal{F}$, respectively, prepare states ρ and σ, we obtain

$$W_{\rho\otimes\sigma} = W_\rho\otimes W_\sigma, \tag{20}$$

which informs us that our chosen representation factorizes over subsystems for product states.

The transition matrix formed by $W_{\mathcal{E}}(\mathbf{v}|\mathbf{u})$ preserves normalization since

$$\sum_{\mathbf{v}\in\mathcal{P}_m} W_{\mathcal{E}}(\mathbf{v}|\mathbf{u}) = 1, \tag{21}$$

for any $\mathbf{u}\in\mathcal{P}_n$ and any quantum channel $\mathcal{E}$. [Proofs of Eqs. (17) through Eq. (21) can be found in Appendix A 3.] Therefore, a quantum channel $\mathcal{E}$ from n qubits to m qubits is represented by a stochastic matrix if and only if $W_{\mathcal{E}}(\mathbf{v}|\mathbf{u})\geq 0$ for all $\mathbf{u},\mathbf{v}$. By inspection of Eq. (16) we equivalently have that the quantum channel $\mathcal{E}$ is stochastically represented if and only if its Choi state $\mathcal{J}(\mathcal{E})$ on $n+m$ qubits is represented by a genuine probability distribution on the phase space $\mathcal{P}_{n+m}$.

Unlike the odd-dimensional case, the channel $\mathcal{E}$ is not guaranteed a stochastic representation whenever $\mathcal{J}(\mathcal{E})$ is a stabilizer state. This is an immediate consequence of the sequential and parallel composition rules of Eqs. (18) and Eq. (19), which imply that our qubit representation cannot be non-negative over the full stabilizer subtheory [53]. However, we show in the next section that $\mathcal{E}$ is stochastically represented if $\mathcal{J}(\mathcal{E})$ belongs to an important subset of stabilizer states known as CSS states.

C. CSS states and circuits

We now identify a class of qubit distillation protocols that arise naturally in fault-tolerant quantum computing, are sufficiently large to enable universal quantum computation, and admit a stochastic representation. In particular, we show that a channel is stochastically represented if its Choi state is CSS. Building on this result, we construct a stochastically represented subset of stabilizer circuits wherein CSS states play the role of stabilizer states, which has been shown to be capable of universal quantum computation with magic state injection. A schematic of our approach is shown in Fig. 2.

A pure CSS state on n qubits is any stabilizer state whose stabilizer group can be generated by n Pauli observables that are individually of X-type or Z-type only. For instance, $|\phi^+\rangle := 1/\sqrt{2}(|00\rangle + |11\rangle)$ has the stabilizer group

$$\mathcal{S}(|\phi^+\rangle) = \langle X_1X_2, Z_1Z_2\rangle, \tag{22}$$

and is therefore CSS. By contrast, $|\psi\rangle := \mathbb{1} \otimes H\,|\phi^+\rangle$ is stabilized by

$$\mathcal{S}(|\psi\rangle) = \langle X_1Z_2, Z_1X_2\rangle, \tag{23}$$

and is not CSS because its stabilizer generators necessarily mix X and Z. As they are generators of stabilizer groups defining CSS states, we group X- and Z-type Pauli observables together as CSS observables. Letting the set of all pure CSS states be denoted Ω_{CSS}, we further define the set of all CSS states $\mathcal{D}_{\text{CSS}}$ as the convex hull of Ω_{CSS}.

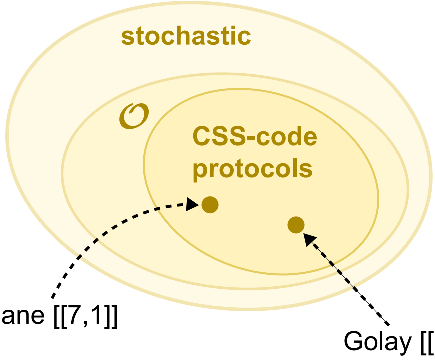

FIG. 2. (Schematic of our approach.) We find that the set of completely CSS-preserving protocols $\mathcal{O}$ are stochastically represented. Such protocols contain the family of CSS-code projections as subset, examples of which include 7-1 and 23-1 protocols based, respectively, on the Steane [[7, 1]] and Golay [[23, 1]] codes [19].

The representation we choose coincides on all rebit states with an earlier one introduced by Ref. [44] (see Appendix A 2), in which it was shown that a Discrete Hudson's theorem [28] can be recovered for qubits when one restricts to rebits. More precisely, it was shown that any pure n-rebit state is non-negatively represented if and only if it is CSS. Therefore, W_ρ is a valid probability distribution for all $\rho \in \mathcal{D}_{\text{CSS}}$, and we conclude the following.

Theorem 1.—A quantum channel $\mathcal{E}$ from n to m qubits is stochastically represented if $\mathcal{J}(\mathcal{E})$ is a CSS state on $n+m$ qubits.

Theorem 1 can be leveraged to identify stochastically represented qubit stabilizer operations in a systematic way. A channel $\mathcal{E}$ from a system of qubits B is CSS preserving if $\mathcal{E}(\rho)$ is a CSS state for all $\rho \in \mathcal{D}_{\text{CSS}}$, and completely CSS preserving if, given any CSS state ρ_{AB} on another system of qubits A as well as B, $\mathcal{I}_A \otimes \mathcal{E}_B(\rho_{AB})$ is always CSS. We now note that the maximally entangled state $|\phi_n^+\rangle$ over two sets of n qubits is CSS for all n (see Appendix A 3). Therefore, if $\mathcal{E}$ is completely CSS preserving, $\mathcal{J}(\mathcal{E})$ must be CSS. By Theorem 1, it follows that every completely CSS-preserving channel is stochastically represented.

To motivate the class of completely CSS preserving channels as operationally significant, we highlight that they cover at least the following subset of stabilizer circuits (see Appendix B 3 for proof).

Lemma 2 (CSS circuits).—Any sequence of the following stabilizer operations:

(1) introducing a CSS state on any number of qubits;
(2) performing a completely CSS-preserving gate on any number n of qubits, i.e., a member of the group

$$\mathcal{G}(n) := \langle \text{CNOT}(i,j), Z_i, X_i\rangle_{i,j=1,\ldots,n, i\neq j}; \tag{24}$$

(3) projectively measuring a CSS observable (with the possibility of classical control conditioned on outcome);
(4) discarding any number of qubits;

as well as statistical mixtures of such sequences, is completely CSS preserving.

By using CSS preserving rather than completely CSS-preserving gates, channels covered by Lemma 2 can be promoted to the subset of stabilizer circuits where CSS states play the role of stabilizer states. However, both groups of gates are equally powerful for magic distillation (see discussion in Appendix B 3 a). Thus we directly refer to the set of channels covered by Lemma 2 as "CSS circuits," and conclude that all such circuits are stochastically represented.

We emphasize the computational power of CSS circuits. Firstly, they are capable of universal quantum computation when supplemented by rebit magic states [44,45] they can also distil [54]. Moreover, the gateset $\mathcal{G}(n)$ constitutes all gates that can be implemented fault tolerantly using defect braiding in surface codes [55]. Finally, we see in Sec. V that CSS circuits form the basis of many existing magic distillation protocols constructed around CSS codes.

IV. ENTROPIC CONSTRAINTS ON COMPLETELY CSS-PRESERVING PROTOCOLS

The standard approach to obtaining constraints on magic distillation is tracking a magic monotone [48,56–59], which is any property of a quantum system that cannot be increased under some class of magic nongenerating operations (e.g., stabilizer operations). The paradigmatic example is mana [56], the total negativity in the Wigner representation of a state. However, this approach operates at the state level and therefore does not incorporate any additional distinguishing physics of distillation protocols. In contrast, the recent work [27] considers how a class of magic distillation protocols transform a *pair* of quantum states—one a noisy magic state, the other a stabilizer state singled out by the characteristic physics of those protocols.

Here we briefly review the approach taken in Ref. [27] to extend relative majorization to quasiprobability distributions, and how that leads to the extension of a dense subset of α-Rényi divergences from classical statistical mechanics to quantify the nonclassical order in magic states under distillation. We then adapt this work for rebit magic state distillation using CSS circuits.

A. Statistical mechanics of quasiprobability distributions

At the heart of statistical mechanics are the notions of disorder and deviations from equilibrium. In classical statistical mechanics, this leads to the thermodynamic entropy $H(\mathbf{p}) = -\sum_i p_i \log p_i$, which is essentially the unique measure of disorder of a statistical distribution $\mathbf{p} = (p_1, \ldots, p_N)$.

In odd-dimensional systems [16] or restricted qubit models [44,45,60], magic states that promote an efficiently simulable part of quantum mechanics to universal quantum computation must have negativities in their representation within a phase-space model. Despite this negativity, it is still possible to arrive at a well-defined statistical mechanical description that circumvents the fact that the Boltzmann entropy is not well defined. The key observation we exploit is that the framework of majorization remains well defined when extended to quasiprobability distributions, and is a more fundamental concept than the traditional entropy.

Given two probability distributions $\mathbf{p} = (p_1, \ldots, p_N)$ and $\mathbf{p}' = (p'_1, \ldots, p'_M)$, we would like to determine which of them is "more disordered" than the other. This can be done by comparing $\mathbf{p}$ to some reference probability distribution $\mathbf{r} = (r_1, \ldots, r_N)$ of our choice, and $\mathbf{p}'$ to some other reference probability distribution $\mathbf{r}' = (r'_1, \ldots, r'_M)$, also of our choice. We then say that $(\mathbf{p}, \mathbf{r})$ relatively majorizes $(\mathbf{p}', \mathbf{r}')$ and write $(\mathbf{p}, \mathbf{r}) \succ (\mathbf{p}', \mathbf{r}')$ if there exists a stochastic map A that sends the first pair of distributions into the second, namely

$$(A\mathbf{p}, A\mathbf{r}) = (\mathbf{p}', \mathbf{r}'). \tag{25}$$

It was shown in Ref. [27] that this definition can be extended to the case of quasiprobability distributions in the first argument, and the following result is established to provide an entropic measure in terms of the α-Rényi divergences.

Theorem 2 ([27]).—Let $\mathbf{w} = (w_1, \ldots, w_N)$ and $\mathbf{w}' = (w'_1, \ldots, w'_M)$ be any two quasiprobability distributions and let $\mathbf{r} = (r_1, \ldots, r_N)$ and $\mathbf{r}' = (r'_1, \ldots, r'_M)$ be any two probability distributions with nonzero components. If $(\mathbf{w}, \mathbf{r}) \succ (\mathbf{w}', \mathbf{r}')$ then

$$D_\alpha(\mathbf{w}||\mathbf{r}) \geq D_\alpha(\mathbf{w}'||\mathbf{r}'), \tag{26}$$

for all $\alpha \in \mathcal{A} = \left\{\frac{2a}{2b-1} \,:\, a, b \in \mathbb{N}, a \geq b\right\} \cup \{\infty\}$.

Here $D_\alpha(\mathbf{w}||\mathbf{r})$ is an extension of the classical α-Rényi divergence to the case of $\mathbf{w}$ being a quasiprobability distribution. This extension requires $\alpha \in \mathcal{A}$ in order for the expression

$$D_\alpha(\mathbf{w}||\mathbf{r}) := \frac{1}{\alpha - 1} \log \sum_{i=1}^{N} w_i^\alpha r_i^{1-\alpha}, \tag{27}$$

to be well defined [61]. In the case of $\mathbf{r}$ being the uniform distribution $\mathbf{r} = (1/N, 1/N, \ldots, 1/N)$, we have that

$$D_\alpha(\mathbf{w}||\mathbf{r}) = \log N - H_\alpha(\mathbf{w}), \tag{28}$$

where $H_\alpha(\mathbf{w}) := (1-\alpha)^{-1} \log \sum_i w_i^\alpha$ is the α-Rényi entropy evaluated on $\mathbf{w}$. Another result of Ref. [27] is that $\mathbf{w}$ has negativity if and only if $H_\alpha(\mathbf{w})$ is negative for α close to 1 and diverges to $-\infty$ in the limit $\alpha \to 1^+$. This provides a well-defined and meaningful notion of negative entropy in a statistical mechanical setting.

B. Application to completely CSS-preserving magic distillation

Since completely CSS-preserving protocols are stochastically represented, the following family of entropic constraints on rebit magic distillation applies them all (see Appendix C for proof).

Theorem 3.—Let ρ be a noisy rebit magic state and τ be a CSS state in the interior of $\mathcal{D}_{\text{CSS}}$. If there exists a completely CSS-preserving protocol $\mathcal{E}$ such that $\mathcal{E}(\rho^{\otimes n}) = \rho'$

and $\tau' := \mathcal{E}(\tau^{\otimes n})$ is also in the interior of $\mathcal{D}_{\text{CSS}}$, then

$$\Delta D_\alpha \geq 0 \tag{29}$$

for all $\alpha \in \mathcal{A}$, where

$$\Delta D_\alpha := nD_\alpha(W_\rho||W_\tau) - D_\alpha(W_{\rho'}||W_{\tau'}). \tag{30}$$

The reference process $\tau^{\otimes n} \mapsto \tau'$ in Theorem 3 can be used in three different ways: (1) as a variational parameter, (2) to account for limitations the physical hardware carrying out magic distillation, or (3) to capture structure distinctive to a family of protocols and thereby produce entropic constraints specific to that family, which we now elaborate on in turn.

For (1), we simply treat τ as a variational parameter, which can be optimized over $\mathcal{D}_{\text{CSS}}$ to obtain the following set of monotones [62] on completely CSS-preserving protocols

$$\Lambda_\alpha(\rho) := \inf_{\tau\in\mathcal{D}_{\text{CSS}}} D_\alpha(W_\rho||W_\tau), \tag{31}$$

for all $\alpha \in \mathcal{A}$. To see this, we note that we have $D_\alpha(W_\rho||W_\tau) \geq 0$ for all ρ, τ, with equality if and only if $\rho = \tau$ (see Lemma 10). Given any rebit state ρ, let τ_ρ be a solution to the optimization problem in Eq. (31). Then if there exists a completely CSS-preserving protocol $\mathcal{E}$ such that $\mathcal{E}(\rho) = \rho'$, we obtain

$$\begin{aligned}\Lambda_\alpha(\rho) = D_\alpha(W_\rho||W_{\tau_\rho}) &\geq D_\alpha(W_{\rho'}||W_{\mathcal{E}(\tau_\rho)})\\ &\geq \Lambda_\alpha(\rho'),\end{aligned} \tag{32}$$

where the first inequality follows from generalized relative majorization and the second inequality follows by the definition in Eq. (31). Therefore $\{\Lambda_\alpha\}_{\alpha\in\mathcal{A}}$ form an infinite set of monotones on all completely CSS-preserving protocols. It is straightforward to verify that Λ_α are subadditive, i.e., $\Lambda_\alpha(\rho^{\otimes n}) \geq n\Lambda_\alpha(\rho)$ (this follows from the additivity of the generalized α-Rényi divergences). Therefore, these Λ_α monotones allow us to set global bounds on any completely CSS-preserving protocol. More precisely, if there exists a completely CSS-preserving protocol $\mathcal{E}$ such that $\mathcal{E}(\rho^{\otimes n}) = \rho'$, then the overhead n is lower bounded as

$$n \geq \frac{\Lambda_\alpha(\rho')}{\Lambda_\alpha(\rho)}. \tag{33}$$

For (2), we can use the reference process to take into account limitations in the hardware carrying out magic distillation. For instance, Ref. [27] uses the reference process to preserve the Gibbs state in order to encode a background temperature or free-energy production in the distillation hardware.

The final way (3) of using the reference process is demonstrated in the next section. We show explicitly how the reference process may be chosen to produce entropic constraints specialized for CSS-code-projection protocols.

V. ENTROPIC CONSTRAINTS ON CSS-CODE-PROJECTION PROTOCOLS

In this section, we apply Theorem 3 to CSS-code-projection protocols, and obtain lower and upper bounds on their code length (which is related to the resource cost). In some parameter regimes, the new lower bounds outperform those due to magic monotones, such as generalized robustness [63] and projective robustness [46]. To our knowledge, these constitute the first set of trade-off relations on distillation parameters that act as fundamental upper bounds on the resource cost for a family of distillation protocols.

A. CSS-code projections

An elementary protocol for magic distillation, proposed in the seminal work of Bravyi and Kitaev [18], uses projection onto a quantum error-correcting code. This protocol begins by taking in n copies of a noisy magic state ρ and postselecting the no-error outcome from the syndrome measurement of an $[[n,k]]$ stabilizer code $\mathcal{C}$. Doing so has the effect of projecting onto $\mathcal{C}$, so the protocol proceeds to decode onto k output qubits and discard the remaining syndrome qubits. In general, any such code-projection protocol succeeds only probabilistically with some acceptance probability p. Nevertheless, if the likelihood of an undetectable error is less than the input error rate ϵ, the postselected output state will have a higher fidelity per qubit with respect to the target magic state than ρ. Many existing magic distillation protocols are based on CSS codes, such as the 15-to-1 protocol [18] based on the $[[15,1]]$ punctured Reed-Muller code [64,65], as well as straightforward code-projection protocols based on the $[[7,1]]$ Steane and $[[23,1]]$ Golay CSS codes analyzed in Ref. [19].

It has long been known that any n-to-1 magic distillation protocol can be decomposed as a sum of stabilizer-code projections followed by Clifford postprocessing [47]. This result implies that the optimal fidelity with respect to a target magic state, though not necessarily optimal acceptance probability, can always be achieved by a stabilizer-code projection. In a similar way, we can show (Theorem 6) that any CSS circuit carrying out an n-to-k magic distillation protocol is a sum of CSS-code projections followed by completely CSS-preserving postprocessing. (In fact, the proof line we give also allows one to generalize the result of Ref. [47] to arbitrary n-k stabilizer protocols).

An n-to-k CSS-code-projection protocol is an operation $\mathcal{K}$ from n to k qubits that acts as

$$\mathcal{K}(\cdot) := \text{tr}_{k+1,\ldots,n}[\mathcal{U} \circ \mathcal{P}(\cdot)], \tag{34}$$

where $\mathcal{U}$ and $\mathcal{P}$ are, respectively, a unitary decoding channel and codespace projection for an $[[n,k]]$ CSS code.

Given n copies of a noisy magic state ρ, $\mathcal{K}$ acts as

$$\mathcal{K}(\rho^{\otimes n}) = p\,\rho', \tag{35}$$

where ρ' is the output magic state on k qubits and we define the acceptance probability $p := \mathrm{tr}[P\rho^{\otimes n}]$ for a single successful run of $\mathcal{K}$. Distillation is successful if the output ρ' from a successful run has a greater fidelity per qubit with respect to a target (pure) magic state of choice than ρ.

Since code-projection protocols are not trace preserving, the majorization constraints do not immediately apply. However, this can be remedied by preparing a specially designated CSS state σ on k qubits whenever an n-to-k code-projection protocol fails, while continuing to distinguish between successful (labeled "0") and unsuccessful (labeled "1") runs of the protocol by recording this information in an ancillary qubit. We can therefore extend $\mathcal{K}$ into the following trace-preserving operation $\mathcal{E}$:

$$\mathcal{E}(\cdot) := \mathcal{K}(\cdot) \otimes |0\rangle\langle 0| + \mathrm{tr}[\overline{\mathcal{P}}(\cdot)]\sigma \otimes |1\rangle\langle 1|, \tag{36}$$

where $\overline{\mathcal{P}} := \overline{P}(\cdot)\overline{P} := (\mathbb{1}^{\otimes n} - P)(\cdot)(\mathbb{1}^{\otimes n} - P)$ performs the projection onto the orthogonal complement of $\mathcal{C}$, and σ is an arbitrary CSS state. We conclude that there exists an n-to-k CSS-code projection such that $\rho^{\otimes n} \mapsto p\,\rho'$ if and only if there exists a trace-preserving n-to-k CSS-code projection $\mathcal{E}$ identified in Eq. (36) such that

$$\mathcal{E}[\rho^{\otimes n}] = p\,\rho' \otimes |0\rangle\langle 0| + (1-p)\sigma \otimes |1\rangle\langle 1| := \rho_p. \tag{37}$$

B. General constraints on CSS-code projections

We now exploit Theorem 3 to derive majorization conditions that apply across *all* n-to-k CSS-code-projection protocols. Crucially, the trace-preserving CSS-code projection identified in Eq. (36) can be implemented as a CSS circuit (see Appendix B 4 for proof), which leads to the following Lemma.

Lemma 3.—Every trace-preserving CSS-code projection can be executed as a sequence of completely CSS preserving operations, and is therefore stochastically represented.

A natural reference process can be chosen for all trace-preserving CSS-code projections by exploiting the fact that their successful components are subunital. To see this, we first note that the identity operator on n qubits can be decomposed as $\mathbb{1}^{\otimes n} = P + \overline{P}$ for the codespace projector P of *any* $[[n,k]]$ CSS code $\mathcal{C}$. The successful component $\mathcal{K}$ in the trace-preserving code projection of $\mathcal{C}$ therefore acts as

$$\mathcal{K}(\mathbb{1}^{\otimes n}) = \mathcal{K}(P + \overline{P}) = \mathcal{K}(P). \tag{38}$$

Since P is the logical identity on k logical qubits, i.e.,

$$P = \sum_{\mathbf{k}\in\{0,1\}^k} |\mathbf{k}_L\rangle\langle\mathbf{k}_L| \equiv \mathbb{1}_L, \tag{39}$$

the decoding of P in Eq. (34) must give an output state that is proportional to the maximally mixed state on k physical qubits, so we obtain $\mathcal{K}(\mathbb{1}^{\otimes n}) \propto \mathbb{1}^{\otimes k}$. This confirms the successful component of any trace-preserving CSS-code projection to be subunital. Furthermore, since P is a rank-2^k projector, the acceptance probability associated with this protocol is $p = \mathrm{tr}\left[P(\mathbb{1}^{\otimes n})/2^n\right] = 2^{k-n}$. Putting everything together, we find that every n-to-k trace-preserving CSS-code projection $\mathcal{E}$ maps the maximally mixed state to

$$\mathcal{E}\left[\left(\frac{\mathbb{1}}{2}\right)^{\otimes n}\right] = \tau_{n,k}, \tag{40}$$

where we define the output state

$$\tau_{n,k} := 2^{k-n}\frac{\mathbb{1}^{\otimes k}}{2^k} \otimes |0\rangle\langle 0| + (1-2^{k-n})\sigma \otimes |1\rangle\langle 1|. \tag{41}$$

Since $\mathbb{1}/2$ and $\tau_{n,k}$ are both full-rank CSS states (for the appropriate choice of $\sigma \in \mathcal{D}_{\mathrm{CSS}}$), Eq. (40) is a valid reference process for *all* trace-preserving CSS-code projections.

We therefore conclude that, if there exists an n-to-k CSS-code projection such that $\rho^{\otimes n} \mapsto p\,\rho'$, then

$$\Delta D_\alpha := nD_\alpha\left(W_\rho \middle|\middle| W_{\frac{1}{2}}\right) - D_\alpha\left(W_{\rho_p}||W_{\tau_{n,k}}\right) \geq 0 \tag{42}$$

for all $\alpha \in \mathcal{A}$. We define ΔD_α over the restricted domain $n \in [k,\infty]$ (as the number of logical qubits cannot exceed the number of physical qubits). We highlight the satisfying fact that ΔD_α is independent of the choice of CSS state σ in Eq. (36). This follows from resource-theoretic arguments as well as properties of the α-Rényi divergence (see Appendix D 2 for details).

C. Bounds on the code length of CSS-code-projection protocols

The entropic constraints of Eq. (42) on CSS-code projections can be used to bound many metrics on their performance as magic distillation protocols. In this section, we apply these constraints to bounding the code length of any CSS-code-projection protocol that could achieve some target combination of noise reduction and acceptance probability from a given supply of noisy magic states.

We first highlight some properties of the relative entropy difference ΔD_α in the following lemma, proofs of which can be found in Appendix D 3.

Lemma 4.—The following properties of the relative entropy difference ΔD_α hold for any noisy input rebit magic state ρ, output k-rebit magic state ρ', acceptance probability $p < 1$ and $\alpha \in \mathcal{A}$:

(i) ΔD_α is concave over the domain $n \in [k,\infty]$.

(ii) ΔD_α is negative in the limit where $n = k$:

$$\lim_{n\to k^+} \Delta D_\alpha < 0. \tag{43}$$

.

(iii) If $H_\alpha[W_\rho] > 1$, then ΔD_α is also negative in the asymptotic limit

$$\lim_{n\to \infty} \Delta D_\alpha < 0. \tag{44}$$

An immediate consequence of Lemma 4 is that ΔD_α is either negative for all n, which implies no CSS-code-projection protocol can distil ρ' from a supply of ρ with acceptance probability p, or ΔD_α has one or two roots located at n_L^α and n_U^α. These roots therefore constitute lower and upper bounds on the code length n of any CSS-code projection that can carry out the desired distillation. We formalize these observations in the following Theorem.

Theorem 4.—Let ρ be a noisy rebit magic state. Any n-to-k CSS-code projection can distil out the k-rebit magic state ρ' from a supply of ρ at acceptance probability p must have a code length n such that

$$n \ge n_L^\alpha := \begin{cases} \inf_n\{n : \Delta D_\alpha \ge 0\} & \text{if } \exists n : \Delta D_\alpha \ge 0, \\ \infty & \text{otherwise.} \end{cases} \tag{45}$$

$$n \le n_U^\alpha := \begin{cases} \sup_n\{n : \Delta D_\alpha \ge 0\} & \text{if } \exists n : \Delta D_\alpha \ge 0, \\ -\infty & \text{otherwise.} \end{cases} \tag{46}$$

for all $\alpha \in \mathcal{A}$. Moreover, given any α such that $H_\alpha[W_\rho] > 1$, the second expression yields a *finite* upper bound on n.

For sufficiently low k, these bounds can be computed numerically using basic root-finding methods. However, we also find analytic upper and lower bounds on n in Sec. V C 1. The values of n_L^α and n_U^α when $D_\alpha < 0$ for all n are chosen to indicate that no n-to-k CSS-code projection can carry out the desired distillation.

We emphasize that n in Theorem 4 refers to the code length (related to the resource cost C by $C = n/pk$) in a single run of a distillation protocol, as opposed to the asymptotic overhead. However, single-run n still constitutes a useful metric for analyzing the actual resource cost of a given stage of a protocol. Moreover, distillation costs are typically dominated by the final round of a multistage distillation protocol (see Ref. [1] and references contained therein), so we expect the above bounds to be particularly informative in this context.

In Fig. 3, we plot the tightest lower bound produced by Theorem 4 on the code length of n-to-1 CSS-code-projection protocols for Hadamard state distillation. We consider input magic states of the form $\rho(\epsilon) := (1-\epsilon)\,|H\rangle\,\langle H| + \epsilon(\mathbb{1})/2$, which can be generally assumed irrespective of the particular error model. This is because any state ρ can be converted into this canonical form by applying the preprocessing channel

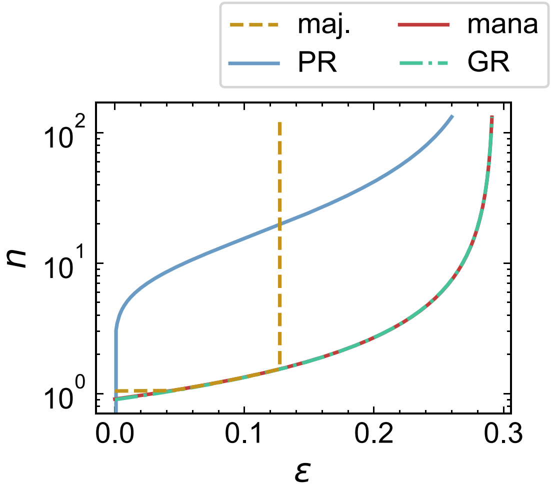

FIG. 3. (Lower bound comparison.) We plot lower bounds on the number of copies n of the noisy Hadamard state $(1-\epsilon)\,|H\rangle\,\langle H| + \epsilon\frac{\mathbb{1}}{2}$ required to distil a single output qubit $|H\rangle$ with output error rate $\delta = 10^{-9}$ and acceptance probability $p = 0.9$ under a CSS-code-projection protocol as a function of input error rate ϵ. Our tightest lower bound from majorization (maj.) is shown to be tighter those from mana [56] and generalized robustness (GR) [63]. However, it only outperforms the lower bound from projective robustness (PR) [46] in the high p, high ϵ regime.

$$\mathcal{E}(\rho) := \frac{1}{2}(\mathbb{1}\rho\mathbb{1} + H\rho H), \tag{47}$$

i.e., twirling with respect to the Clifford subgroup generated by the single Hadamard gate. In all parameter regimes, our lower bound is observed to be tighter than mana [66], and the generalized robustness bound in Theorem 13 of Ref. [63]. Furthermore, in the high p, high ϵ regime our lower bound gives tighter constraints than the projective robustness bound [46]. In particular, there is a cutoff input error rate $\epsilon \approx 0.12$ at which our lower bound shoots up to infinity because, for any input error greater than this cutoff, one can always find some α such that $\Delta D_\alpha < 0$ for all n, so no CSS-code projection can carry out the desired distillation given a higher input error rate (see Sec. V C 2 for physical intuition on the origin of this behavior). In the low-p regime, our *upper bounds* are still able to give additional constraints on code length beyond those given by the projective robustness bound. In particular, Fig. 1 puts together information from our upper bounds and the lower bound from projective robustness to show that no CSS-code projection can achieve some target combinations of output error and acceptance probability beyond a cutoff input error rate.

1. Extension to nonqubit code projection

Mathematically, the appearance of upper bounds on the code length of qubit CSS-code-projection protocols comes from the concavity of the objective function ΔD_α in n. We now demonstrate that this feature is not peculiar to qubits, and in fact arises whenever stabilizer code projections on any quantum system are stochastic under a Wigner representation sufficiently similar to Gross's.

More precisely, we say that a Wigner representation W of n qudits with Hilbert-space dimension d is a *generalised Gross's Wigner representation* if it represents each state ρ as the function W_ρ on a phase space $\mathbb{Z}_d^n \times \mathbb{Z}_d^n$,

$$W_\rho(\mathbf{u}) := \frac{1}{d^n}\mathrm{tr}\left[A_{\mathbf{u}}^\dagger \rho\right], \tag{48}$$

via phase-point operators $\{A_{\mathbf{u}}\}$ satisfying (A1)–(A4). For all stabilizer-code projections on odd-dimensional systems, W can simply be Gross's Wigner representation [27]. For CSS-code projections on qubits, W can be the representation from Eq. (8). We then have the following analytic upper and lower bounds on the resource cost of code-projection protocols.

Theorem 5 (Qudit-code bounds).—Consider the distillation of k copies of a pure magic state ψ from a supply of the noisy magic state ρ, where ψ and ρ are d-dimensional qudit states that are real represented under a generalized Gross's Wigner representation W. Any stochastically represented distillation protocol that projects onto the codespace of an $[[n,k]]$ stabilizer code and can use n copies of ρ to distil out a k-qudit state ρ' with acceptance probability p and output error $\delta \geq \|\rho' - \psi^{\otimes k}\|_1$ must have a code length n such that

$$n \geq \frac{k\left[\log d - H_\alpha(W_\psi)\right] - \frac{\alpha}{1-\alpha}\log\left(\frac{p}{1+\delta d^{5/2}}\right)}{\left[\log d - H_\alpha(W_\rho)\right]}, \tag{49}$$

for all $\alpha \in \mathcal{A}$ for which $H_\alpha(W_\rho) < \log d$, and

$$n \leq \frac{k\left[H_\alpha(W_\psi) - \log d\right] + \frac{\alpha}{1-\alpha}\log\left(\frac{p}{1+\delta d^{5/2}}\right)}{\left[H_\alpha(W_\rho) - \log d\right]}, \tag{50}$$

for all $\alpha \in \mathcal{A}$ for which $H_\alpha(W_\rho) > \log d$.

One might be concerned that the conditions $H_\alpha(W_{\rho(\epsilon)}) > \log d$ given in Theorems 4 and 5 for the existence of a finite upper bound on n are never actually satisfied. This turns out not to be the case. For n-to-1 CSS-code-projection protocols for Hadamard distillation, there always exists a valid set of α values such that $H_\alpha[W_{\rho(\epsilon)}] > 1$ for all ϵ, so we always obtain a finite upper bound on n. This can be seen by upper bounding the Rényi entropy of $W_{\rho(\epsilon)}$ as

$$H_\alpha\left[W_{\rho(\epsilon)}\right] \geq H_\alpha\left[W_{\rho(0)}\right], \tag{51}$$

and then examining Fig. 4, which shows that there exists a finite range of α such that $H_\alpha > 1$ at $\epsilon = 0$.

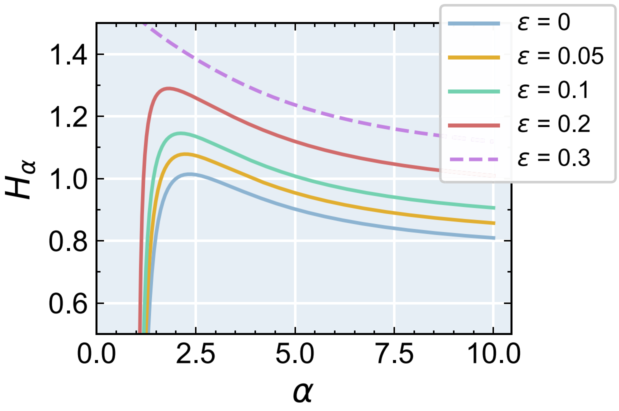

FIG. 4. (Wigner-Rényi entropies and magic distillation.) We plot the condition in Theorem 4 for the existence of finite upper bounds on n in an n-to-1 CSS distillation for the qubit Hadamard state distillation, signified by region where $H_\alpha[W_{\rho(\epsilon)}] > 1$. Even in the limit of zero input error $\epsilon = 0$ we obtain a valid set of permissible α, which implies that Hadamard state distillation under n-to-1 CSS-code projection is ruled out in the asymptotic limit $n \to \infty$. We further highlight that the error rate $\epsilon = 0.3$ (dashed curve) is outside of the region where $\rho(\epsilon)$ is magic ($0 \leq \epsilon < 1 - \frac{1}{\sqrt{2}}$), and therefore $W_{\rho(\epsilon)}$ is a proper probability distribution at $\epsilon = 0.3$, which is why H_α is only seen to satisfy standard monotonicity properties at this input error. We also highlight that the $\alpha \to 1$ divergence corresponds to a pole in $H_\alpha[W_\rho]$ for magic state ρ, and its residue is the mana of the state.

The trade-off relations given in Theorem 5 can be rearranged to bound other parameters such as the acceptance probability p. In Fig. 5, we compare the tightest upper bound on p produced by Theorem 5 for n-to-1 Hadamard distillation to those attained in existing protocols based on CSS codes given in Refs. [18] and [19]. We see that the acceptance probabilities of basic code-projection protocols using the $[[7,1]]$ Steane and $[[23,1]]$ Golay codes in Fig. 5(a) are orders of magnitude less than our upper bounds, suggesting that substantial room for improvement is not ruled out. Interestingly, in Fig. 5(b) our upper bound is very close to the actual acceptance probability of the protocol based on the $[[15,1]]$ Reed-Muller code in Ref. [18], which we speculate may hint at something fundamental about the role of the intermediate Clifford corrections used in that protocol.

Theorem 5 takes on a particularly simple form in the particular case of n-to-1 CSS-code-projection protocols for Hadamard-state distillation. By evaluating the $\alpha = 2$ condition explicitly, we find that, if there exists a CSS-code projection that sends n copies of $\rho(\epsilon)$ to a k-qubit state ρ' with acceptance probability p and output error $\delta \geq \|\rho' - |H\rangle\langle H|^{\otimes k}\|_1$, then n is upper bounded as

$$n \leq n^* := 2\log_{f(\epsilon)}\left[\frac{1+6\delta}{p}\right], \tag{52}$$

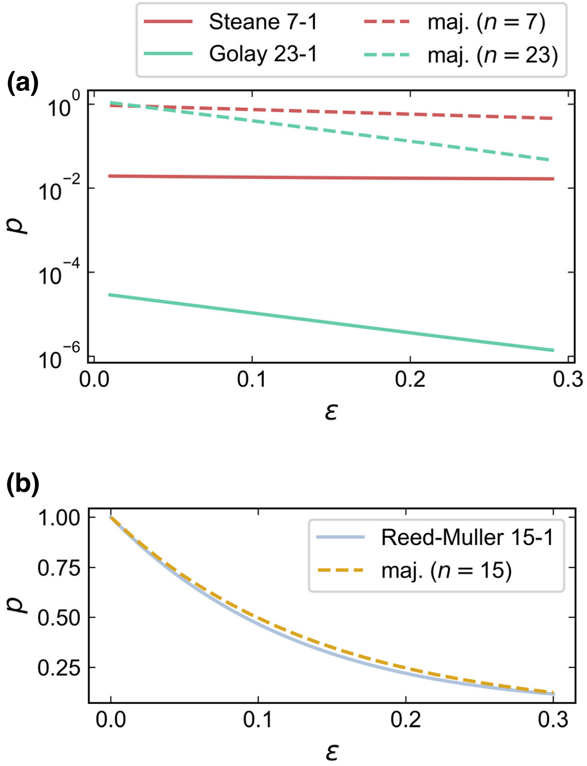

FIG. 5. (Explicit protocol comparison.) (a) We compare the majorization upper bounds (dashed lines) on the acceptance probability p with which one can distil a noisy Hadamard state $(1-\epsilon)\,|H\rangle\langle H| + \epsilon\frac{\mathbb{1}}{2}$ via an n-to-1 code projection against actual acceptance probabilities attained using the Steane code (purple) at $n = 7$ and the Golay code (green) at $n = 23$ (detailed in Ref. [19]). Attained acceptance probabilities are orders of magnitude less than our upper bounds. (b) We plot the majorization upper bound (dashed line) on the acceptance probability p of any 15-to-1 CSS-code-projection protocol with which one can distil the noisy magic state $(1-\epsilon)\,|A\rangle\langle A| + \epsilon\frac{\mathbb{1}}{2}$. Interestingly, our bound is very close to the actual acceptance probability for the 15-to-1 protocol (blue line) given in Ref. [18], though we emphasize this latter protocol is *not* a straightforward CSS-code projection.

where the logarithm base is $f(\epsilon) := [1-\epsilon+\epsilon^2/2]^{-1}$. This expression captures the fact that under a CSS-code-projection protocol, there is a fundamental trade-off between acceptance probability and output fidelity. For instance, Eq. (52) shows that given a supply of noisy magic states ($\epsilon > 0$), we cannot use CSS-code projection to distil a perfect magic state ($\delta = 0$) with certainty ($p = 1$), which was first shown in Ref. [48]. To further investigate this trade-off, in Fig. 6(b) we plot the *maximum achievable fidelity* with respect to the Hadamard state that can be achieved by an n-to-1 CSS-code projection $\mathcal{K}$ via

$$F_{\max}(\rho) = \max_{\mathcal{K}} \left\{ \langle H|\,\rho'\,|H\rangle \;:\; \mathcal{K}(\rho^{\otimes n}) \mapsto p\,\rho' \right\}, \qquad (53)$$

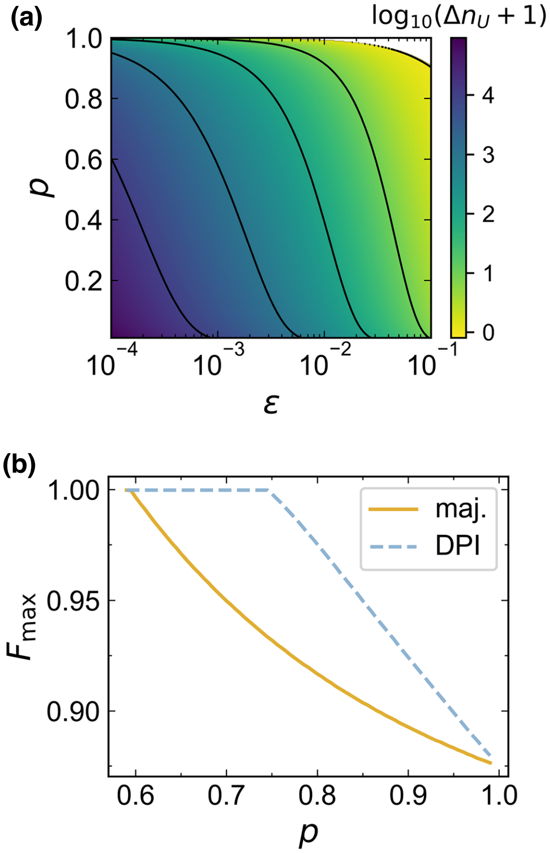

FIG. 6. (Majorization gives independent constraints over DPI.) (a) Shown is how (scaled) $\Delta n_U := n_U^{\mathrm{DPI}} - n_U^{\mathrm{maj}}$ varies over all possible values of acceptance probability p and a realistic range of input error ϵ, with fixed $\delta = 10^{-9}$. Whenever we have $\log_{10}(\Delta n_U + 1) > 0$ means that upper bounds from majorization give tighter constraints than the DPI, reaching $\Delta n_U = O(10^4)$ in the low p, low ϵ regime. (b) We show the trade-off relation given by bounds on the maximum achievable fidelity $F_{\max}(\rho)$ versus target acceptance probability p, under an n-1 CSS-code projection, where $\rho = \frac{3}{4}\,|H\rangle\langle H| + \frac{1}{8}\mathbb{1}$. For p above a given threshold (approximately equal to 0.6) no perfect distillation is theoretically possible, even for $n \to \infty$ copies of the input state. Majorization (maj.) is shown to give stronger constraints than that of DPI.

where the maximization is performed over the set of all n-to-1 CSS-code-projection protocols.

2. *Why do we expect upper bounds?*

Taking CSS circuits as our free operations, the appearance of upper bounds on n might first seem to contradict a resource theory perspective, where we might expect $n+1$ copies of a noisy magic state to be at least as good as n copies at distilling magic, since discarding subsystems is itself a CSS circuit. However, by specializing to stabilizer-code-projection protocols, we necessarily introduce a trade-off between n and acceptance probability p, which we now show in a simple calculation. For any n-to-k

stabilizer-code-projection protocol for d-dimensional qudit distillation, the acceptance probability p is given by how much of n copies of the noisy input magic state ρ projects onto the d^k-dimensional codespace spanned by the logical basis $\{|j_L\rangle\}_{j=0,\ldots,d^k-1}$ of the code. Letting $\lambda_{\max}(\cdot)$ denote a state's largest eigenvalue, we immediately identify the following upper bound on p:

$$p = \sum_{j=0}^{d^k-1} \langle j_L|\, \rho^{\otimes n}\, |j_L\rangle \\ \leq d^k \lambda_{\max}(\rho^{\otimes n}) = d^k[\lambda_{\max}(\rho)]^n, \tag{54}$$

which falls monotonically towards 0 as $n \to \infty$. At an intuitive level, the trade-off between n and p occurs because the codespaces of $[[n,k]]$ stabilizer codes remain the same size as we increase n, and so take up a vanishingly small part in the support of all the noisy input magic states used. Under the requirement that we have some threshold acceptance probability (below which the expected overhead would be too large), a corresponding upper bound on n is then expected.

3. Comparison with the data-processing inequality (DPI)

We see that stochastically represented CSS-code-projection protocols give rise to a set of upper bounds on n (Theorem 5). By comparing to the DPI, we see that although the existence of upper bounds is a general feature of code-projection protocols, exploiting the stochasticity in the representation of CSS-code projections gives strictly stronger bounds.

According to the DPI, if there exists a code projection (CSS or otherwise) that can distil out the k-qubit magic state ρ' from n copies of the noisy magic state ρ with acceptance probability p, then

$$\Delta\tilde{D}_\alpha := n\tilde{D}_\alpha\left(\rho\middle|\middle|\frac{\mathbb{1}}{2}\right) - \tilde{D}_\alpha\left(\rho_p||\tau_{n,k}\right) \geq 0, \tag{55}$$

for all $\alpha \in (1,\infty)$ [67], where $\tilde{D}_\alpha(\rho||\tau)$ is the sandwiched α-Rényi divergence [68,69] on the normalized quantum states ρ and τ, which is defined as

$$\tilde{D}_\alpha(\rho||\tau) := \frac{1}{\alpha-1}\log \mathrm{tr}\left[\left(\tau^{\frac{1-\alpha}{2\alpha}}\rho\tau^{\frac{1-\alpha}{2\alpha}}\right)^\alpha\right]. \tag{56}$$

We then have the following upper bound on n:

$$n \leq n_U^{\mathrm{DPI}} := \min_{\alpha\in(1,\infty)} \max_n\{n : \Delta\tilde{D}_\alpha \geq 0\}, \tag{57}$$

which is finite whenever α is such that $H_\alpha(\rho) > 1$ (proof essentially identical to that of Theorem 4).

Given that we also obtain upper bounds on code length from simple data processing of quantum states, we should ask, does majorization give genuinely new constraints on magic state distillation beyond the DPI? Since our majorization conditions are a consequence of the stochastic representation of stabilizer circuits, while the DPI arises from the fact that all quantum channels are CPTP, this question may be loosely rephrased as asking whether stochasticity imposes any additional constraints beyond those imposed by CPTP on magic state distillation by stabilizer-code projection.

Figure 6 allows us to answer this question in the affirmative, since the upper bound on code length due to majorization is observed to be stronger than that due to the DPI over a wide range of parameter regimes for CSS-code-projection protocols. In particular, extracting $n_U^{\mathrm{maj}} := \min_\alpha n_U^\alpha$ as the tightest bound due to majorization from Theorem 4, we find that, in the low acceptance probability p and low input error ϵ regime of Fig. 6(a), the difference in upper bounds $\Delta n_U := n_U^{\mathrm{DPI}} - n_U^{\mathrm{maj}}$ (the amount by which majorization "beats" DPI) is of the order $\Delta n_U = O(10^4)$. We thus conclude that the constraints on CSS protocols stemming from majorization go beyond those from the DPI.

VI. BEYOND REBIT DISTILLATION

We simplify our analysis thus far by restricting our attention to rebit magic states such as $|H\rangle$. However, many protocols such as the seminal 15-to-1 Bravyi-Kitaev protocol [18] distil magic states such as $|A\rangle \propto |0\rangle + e^{i(\pi)/4}|1\rangle$, which have complex density matrices in the computational basis. An argument can be made that since $|A\rangle$ is Clifford equivalent to $|H\rangle$, this state should be considered equally resourceful for magic distillation. However, to address this concern more directly, we extend our majorization relations to states with complex Wigner representations. A discussion of how this can be achieved is given in Appendix F. The basic idea is that we can define valid $2d^2$-dimensional quasiprobability distributions by forming the direct sum of the real and imaginary parts of the original distribution:

$$\mathbf{w}_\rho := \mathrm{Re}W_\rho \oplus \mathrm{Im}W_\rho, \tag{58}$$

while for our reference CSS states we can simply take

$$\mathbf{r}_\tau := \frac{1}{2}W_\tau \oplus W_\tau. \tag{59}$$

It then follows that if there exists a completely CSS-preserving operation $\mathcal{E}$ such that $\mathcal{E}(\rho) = \rho'$ and $\mathcal{E}(\tau) = \tau'$, then

$$D_\alpha(\mathbf{w}_\rho||\mathbf{r}_\tau) \geq D_\alpha(\mathbf{w}_{\rho'}||\mathbf{r}_{\tau'}), \tag{60}$$

for all $\alpha \in \mathcal{A}$. Further study of the significance of these complex relative majorization conditions may be of foundational interest, but we leave this for future work.

VII. DISCUSSION AND OUTLOOK

We have shown that the statistical mechanical framework of Ref. [27] can be extended to the experimentally significant case of qubit systems by focusing on the processing of magic states under CSS circuits—i.e., the subset of stabilizer circuits where CSS states play the role of stabilizer states. To achieve this, we made use of a Wigner representation first introduced in Ref. [43] wherein completely CSS-preserving channels correspond to stochastic transformations on phase space. This set of channels include CSS circuits, which are sufficient for universal quantum computing [44,45] and consist precisely of the gateset that can be performed fault-tolerantly on surface-code constructions [70].

Within this framework, we show that relative majorization can be used to encode particular properties of an important class of distillation protocols that project onto CSS codes, in terms of which all protocols carried out by CSS circuits can be decomposed. We establish general entropic constraints on such protocols in terms of upper and lower bounds on the code length n.

In the context of achieving full fault tolerance, a natural extension of this work would be to generalize our results to more sophisticated protocols. For instance, we might ask how the use of m intermediate Clifford corrections in between the measurements of stabilizer generators might affect these fundamental constraints. One would expect to be able to obtain more refined bounds as a function of m. Moreover, while many protocols are based on CSS codes in part due to their relative ease of construction via triorthogonal matrices [20], from an operational perspective it would be of interest to see whether we can extend to the complete set of stabilizer operations on qubit systems.

We have also obtained a set of monotones $\{\Lambda_\alpha\}$ for completely CSS-preserving magic distillation, each of which forms a convex optimization problem. We speculate that an analogous monotone can be constructed for any resource theory for which the free operations are a subset of operations that completely preserve Wigner positivity. From the perspective of quantum optics experiments, wherein Gaussian operations and probabilistic randomness are readily available, it may be of interest to consider the case of continuous variable systems under the set of Gaussian operations and statistical mixtures [71]. Since the individual α-Rényi divergences on quasiprobability distributions were seen in Sec. V C 3 to typically produce stronger constraints than the corresponding constraints given by α-Rényi divergences on quantum states, it would be interesting to see how well these quasidistribution-based monotones perform relative to known state-based counterparts.

On a technical note regarding majorization theory, we point to two interesting directions for further study. Firstly, complex majorization constraints arise naturally when we extend our setup from rebit to all qubit states, where Wigner representations can become complex due to the non-Hermiticity of the operator basis $\{A_{\mathbf{u}}\}$. We expect such constraints to take the form of a duplet of constraints applying separately to the real parts and imaginary parts of the Wigner function. In the context of non-Hermitian quantum mechanics [72], results on complex majorization would also benefit theories that require an ordering between Hamiltonian eigenvalues, such as quantum thermodynamics. Secondly, the Wigner representation of Ref. [43] recovers the covariance over symplectic affine transformations on qubit phase spaces, a property shared by Gross's Wigner function on odd-dimensional systems. This added structure on the phase space was ignored by our analysis, but could be utilized to tighten the obtained bounds in future work. In particular, as explained in the discussion of Ref. [27], the stochastic majorization used in our analysis is only a special case of G majorization, where G can be taken as a subgroup of the stochastic group such as the symplectic group. It can then be shown [73–75] that one should expect to obtain a set of finite lower-bound constraints on distillation, which will be tighter than stochastic majorization constraints.

ACKNOWLEDGMENTS

R.A. and N.K. are supported by the EPSRC Centre for Doctoral Training in Controlled Quantum Dynamics. S.G.C. is supported by the Bell Burnell Graduate Scholarship Fund. D.J. is supported by the Royal Society and a University Academic Fellowship.

APPENDIX A: WIGNER REPRESENTATION

For any n-qubit state ρ, we define the following complex-valued function $W_\rho : \mathcal{P}_n \mapsto \mathbb{C}$ as

$$W_\rho(\mathbf{u}) := \frac{1}{2^n}\mathrm{tr}[A_{\mathbf{u}}^\dagger \rho], \tag{A1}$$

where $\{A_{\mathbf{u}}\}$ are the set of 2^{2n} *phase-point operators* on n-qubits, which are defined as follows:

$$A_{\mathbf{0}} := \frac{1}{2^n}\sum_{\mathbf{u}\in\mathcal{P}_n} D_{\mathbf{u}}, \quad A_{\mathbf{u}} := D_{\mathbf{u}} A_{\mathbf{0}} D_{\mathbf{u}}^\dagger. \tag{A2}$$

As a consequence of Eq. (7), these phase-point operators can alternatively be expressed as

$$A_{\mathbf{u}} = \frac{1}{2^n}\sum_{\mathbf{v}\in\mathcal{P}_n} (-1)^{[\mathbf{u},\mathbf{v}]} D_{\mathbf{v}}, \tag{A3}$$

which further reveals that every $A_{\mathbf{u}}$ is real in the computational basis.

Despite being complex valued, W_ρ transforms covariantly under the displacement operators—informally speaking, ρ is shifted by the displacement operators around phase space—just like Gross' representation in odd dimensions. Concretely, we consider the Wigner representation of $D_{\mathbf{v}}\rho D_{\mathbf{v}}^{\dagger}$ for an arbitrary phase-space displacement $\mathbf{v}$, which is

$$W_{D_{\mathbf{v}}\rho D_{\mathbf{v}}^{\dagger}}(\mathbf{u}) = \frac{1}{2^n}\mathrm{tr}[A_{\mathbf{u}}^{\dagger}D_{\mathbf{v}}\rho D_{\mathbf{v}}^{\dagger}] = \frac{1}{2^n}\mathrm{tr}[D_{\mathbf{v}}^{\dagger}A_{\mathbf{u}}^{\dagger}D_{\mathbf{v}}\rho]$$
$$= \frac{1}{2^n}\mathrm{tr}\left[\sum_{\mathbf{a}\in\mathcal{P}_n}(-1)^{[\mathbf{u},\mathbf{a}]}(D_{\mathbf{a}}D_{\mathbf{v}})^{\dagger}D_{\mathbf{v}}\rho\right]. \quad \text{(A4)}$$

Inserting the commutation relation for the displacement operators in Eq. (7) into Eq. (A4), we obtain

$$W_{D_{\mathbf{v}}\rho D_{\mathbf{v}}^{\dagger}}(\mathbf{u}) = \frac{1}{2^n}\mathrm{tr}\left[\left(\sum_{\mathbf{a}\in\mathcal{P}_n}(-1)^{[\mathbf{u}+\mathbf{v},\mathbf{a}]}D_{\mathbf{a}}^{\dagger}\right)\rho\right]$$
$$= W_{\rho}(\mathbf{u}+\mathbf{v}), \quad \text{(A5)}$$

which confirms that W_ρ transforms covariantly under the action of the displacement operators.

1. Properties of qubit phase-point operators and Wigner function

We first establish that the phase-point operators of a joint system are simply tensor products of phase-point operators on its subsystems:

(A1) *(Factorization).* On a bipartite system with subsystems X and Y, $A_{\mathbf{u}_X\oplus\mathbf{u}_Y} = A_{\mathbf{u}_X}\otimes A_{\mathbf{u}_Y}$.

Proof.—From the definition of $D_{\mathbf{u}}$ in Eq. (6), it is clear that

$$D_{\mathbf{u}} = D_{\mathbf{u}_X}\otimes D_{\mathbf{u}_Y}. \quad \text{(A6)}$$

Let n_X and n_Y be the numbers of qubits in subsystems X and Y, respectively. Then the zero phase-point operator on the bipartite system, $A_{\mathbf{0}}$, is

$$A_{\mathbf{0}} := \frac{1}{2^{n_X+n_Y}}\sum_{\mathbf{u}\in\mathcal{P}_{XY}} D_{\mathbf{u}}$$
$$= \frac{1}{2^{n_X+n_Y}}\sum_{\mathbf{u}_X\in\mathcal{P}_X,\mathbf{u}_Y\in\mathcal{P}_Y} D_{\mathbf{u}_X\oplus\mathbf{u}_Y}$$
$$= \frac{1}{2^{n_X+n_Y}}\sum_{\mathbf{u}_X\in\mathcal{P}_X}\sum_{\mathbf{u}_Y\in\mathcal{P}_Y} D_{\mathbf{u}_X}\otimes D_{\mathbf{u}_Y}$$
$$= A_{\mathbf{0}_X}\otimes A_{\mathbf{0}_Y}, \quad \text{(A7)}$$

which in turn implies that any phase-point operator $A_{\mathbf{u}} := A_{\mathbf{u}_X\oplus\mathbf{u}_Y}$ for some $\mathbf{u}_X\in\mathcal{P}_X$ and $\mathbf{u}_Y\in\mathcal{P}_Y$ is

$$A_{\mathbf{u}} := A_{\mathbf{u}_X\oplus\mathbf{u}_Y} = D_{\mathbf{u}_X\oplus\mathbf{u}_Y}A_{\mathbf{0}}D_{\mathbf{u}_X\oplus\mathbf{u}_Y}^{\dagger}$$
$$= \left(D_{\mathbf{u}_X}A_{\mathbf{0}_X}D_{\mathbf{u}_X}^{\dagger}\right)\otimes\left(D_{\mathbf{u}_Y}A_{\mathbf{0}_Y}D_{\mathbf{u}_Y}^{\dagger}\right)$$
$$= A_{\mathbf{u}_X}\otimes A_{\mathbf{u}_Y}, \quad \text{(A8)}$$

as claimed. ■

Property (A1) enables us to break down any n-qubit phase-point operator $A_{\mathbf{u}}$ as a tensor product of single-qubit phase-point operators,

$$A_{\mathbf{u}} = \bigotimes_{i=1}^{n} A_{\mathbf{u}_j}, \quad \mathbf{u} = \bigoplus_{i=1}^{n}\mathbf{u}_j, \quad \text{(A9)}$$

where $\mathbf{u}_j\in\mathbb{Z}_2\times\mathbb{Z}_2$ is a co-ordinate in the phase space of the j th qubit *only*. It is therefore instructive to calculate the single-qubit phase point operators, which are

$$A_{0,0} = \frac{1}{2}(\mathbb{1}+X+Z+iY),$$
$$A_{0,1} = \frac{1}{2}(\mathbb{1}-X+Z-iY),$$
$$A_{1,0} = \frac{1}{2}(\mathbb{1}+X-Z-iY),$$
$$A_{1,1} = \frac{1}{2}(\mathbb{1}-X-Z+iY). \quad \text{(A10)}$$

We next demonstrate how the explicit forms of single-qubit phase-point operators can be leveraged via Eq. (A9) to prove two further properties for general n-qubit phase point operators. In particular, we show how distinct n-qubit phase-point operators are orthogonal under the Hilbert-Schmidt inner product:

(A2) *(Orthogonality).* Let $A_{\mathbf{u}}$ and $A_{\mathbf{v}}$ be two n-qubit phase-point operators. Then $\mathrm{tr}[A_{\mathbf{u}}^{\dagger}A_{\mathbf{v}}] = 2^n\delta_{\mathbf{u},\mathbf{v}}$.

Proof.—Let us first decompose $\mathbf{u}$ and $\mathbf{v}$ as $\mathbf{u} = \bigoplus_{i=1}^{n}\mathbf{u}_j$ and $\mathbf{v} = \bigoplus_{i=1}^{n}\mathbf{v}_j$, where $\mathbf{u}_j$ and $\mathbf{v}_j$ are phase point co-ordinates on the j th qubit only. By Eq. (A10),

$$\mathrm{tr}[A_{\mathbf{u}_j}^{\dagger}A_{\mathbf{v}_j}] = 2\delta_{\mathbf{u}_j,\mathbf{v}_j}. \quad \text{(A11)}$$

Therefore,

$$\mathrm{tr}[A_{\mathbf{v}}^{\dagger}A_{\mathbf{u}}] = \prod_{j=1}^{n}\mathrm{tr}[A_{\mathbf{u}_j}^{\dagger}A_{\mathbf{v}_j}] = \prod_{j=1}^{n}2\delta_{\mathbf{v}_j,\mathbf{u}_j}$$
$$= 2^n\delta_{\mathbf{u},\mathbf{v}}. \quad \text{(A12)}$$

as claimed. ■

There are $|P_n| = |\mathbb{Z}_2^n \times \mathbb{Z}_2^n| = 4^n$ phase-point operators on n qubits. Property (A2) thus implies $\{A_{\mathbf{u}}\}_{\mathbf{u}\in\mathcal{P}_n}$ forms an orthogonal complex basis for the complex vector space $\mathbb{C}^{2^n\times 2^n}$ of $2n \times 2n$ complex matrices under the Hilbert-Schmidt inner product. Therefore, W_ρ is an *informationally complete* representation of n-qubit states. More precisely, any n-qubit quantum state ρ can be uniquely decomposed as

$$\rho = \sum_{\mathbf{u}} W_\rho(\mathbf{u}) A_{\mathbf{u}}, \tag{A13}$$

where $W_\rho(\mathbf{u}) := 1/(2^n)\mathrm{tr}[A_{\mathbf{u}}^\dagger \rho]$ is a complex function on $\mathcal{P}_n$.

Furthermore, every phase-point operator has trace 1:

(A3) *(Unit trace)*. Let $A_{\mathbf{u}}$ be any n-qubit phase point operator. Then we have $\mathrm{tr}[A_{\mathbf{u}}] = 1$.

Proof.—Let us first decompose $\mathbf{u}$ as $\mathbf{u} = \bigoplus_{j=1}^n \mathbf{u}_j$, where $\mathbf{u}_j$ is a point in the phase space of the jth qubit. We see that $\mathrm{tr}[A_{\mathbf{u}_j}] = 1$ from Eq. (A10). Therefore,

$$\mathrm{tr}[A_{\mathbf{u}}] = \prod_{j=1}^n \mathrm{tr}[A_{\mathbf{u}_j}] = \prod_{j=1}^n 1 = 1, \tag{A14}$$

which completes the proof. ■

Property (A3) implies that all n-qubit functions are *normalized*. Since any n-qubit state ρ has trace 1,

$$\begin{aligned} 1 = \mathrm{tr}[\rho] &= \mathrm{tr}\left[\sum_{\mathbf{u}\in\mathcal{P}_n} W_\rho(\mathbf{u}) A_{\mathbf{u}}\right] \\ &= \sum_{\mathbf{u}\in\mathcal{P}_n} W_\rho(\mathbf{u}) \mathrm{tr}[A_{\mathbf{u}}] \\ &= \sum_{\mathbf{u}\in\mathcal{P}_n} W_\rho(\mathbf{u}), \end{aligned} \tag{A15}$$

where the last equality is established by the unit trace of phase-point operators.

We also find it useful to identify the following property of phase-point operators:

(A4) *(Completeness relation)*.

$$\sum_{\mathbf{u}\in\mathcal{P}_n} A_{\mathbf{u}} = 2^n \mathbb{1}_n \tag{A16}$$

Proof.—Adopting the decomposition of each $A_{\mathbf{u}}$ in Eq. (A9), we see that

$$\begin{aligned} \sum_{\mathbf{u}\in\mathcal{P}_n} A_{\mathbf{u}} &= \sum_{\mathbf{u}_1\in\mathcal{P}_1}, \ldots, \sum_{\mathbf{u}_n\in\mathcal{P}_1} \left(\bigotimes_{j=1}^n A_{\mathbf{u}_j}\right) \\ &= \bigotimes_{j=1}^n \left(\sum_{\mathbf{u}_j\in\mathcal{P}_1} A_{\mathbf{u}_j}\right). \end{aligned} \tag{A17}$$

Using the explicit forms of single-qubit phase-point operators in Eq. (A10), we calculate that

$$\sum_{\mathbf{v}\in\mathcal{P}_1} A_{\mathbf{v}} = 2\mathbb{1}, \tag{A18}$$

from which the result immediately follows. ■

2. Wigner representation for rebits

Any n-qubit state ρ can be decomposed as

$$\rho = \left[\frac{1}{2}\left(\rho + \rho^T\right)\right] + i\left[\frac{-i}{2}\left(\rho - \rho^T\right)\right], \tag{A19}$$

where the transposition is taken with respect to the computational basis. Because $\rho^* = \rho^T$ in any basis, we can identify

$$\rho^{(0)} := \frac{1}{2}(\rho + \rho^T) = \mathrm{Re}(\rho), \tag{A20}$$

$$\rho^{(1)} = i\left[\frac{-i}{2}(\rho - \rho^T)\right] = \mathrm{Im}(\rho), \tag{A21}$$

i.e., $\rho^{(0)}$ and $\rho^{(1)}$ are, respectively, the real and imaginary components of the density matrix of ρ in the computational basis.

We first prove Lemma 1, which shows there is a direct correspondence between the real and imaginary components of a state's Wigner representation and those of its density matrix in the computational basis.

Lemma 1.—Given any n-qubit quantum state ρ,

$$\mathrm{Re}[W_\rho(\mathbf{u})] = W_{\mathrm{Re}(\rho)}(\mathbf{u}) \tag{12}$$

$$\mathrm{Im}[W_\rho(\mathbf{u})] = W_{\mathrm{Im}(\rho)}(\mathbf{u}) \tag{13}$$

for all $\mathbf{u} \in \mathcal{P}_n$, where $\mathrm{Re}(\rho)$ and $\mathrm{Im}(\rho)$ are, respectively, the real and imaginary parts of the density matrix of ρ in the computational basis.

Proof.—Adopting the identification $\rho^{(0)} = \mathrm{Re}(\rho)$ and $\rho^{(1)} = \mathrm{Im}(\rho)$, we can then decompose $W_\rho(\mathbf{u})$ as

$$W_\rho(\mathbf{u}) = \frac{1}{2^n}\mathrm{tr}\left[A_{\mathbf{u}}^\dagger \rho^{(0)}\right] + i\frac{1}{2^n}\mathrm{tr}\left[A_{\mathbf{u}}^\dagger \rho^{(1)}\right]. \tag{A22}$$

Since $A_{\mathbf{u}}$ is real, and both $\rho^{(0)}$ and $\rho^{(1)}$ are real by construction, we conclude that $\mathrm{tr}\left[A_{\mathbf{u}}^\dagger \rho^{(0)}\right]$ and $\mathrm{tr}\left[A_{\mathbf{u}}^\dagger \rho^{(1)}\right]$ are both real for all $\mathbf{u} \in \mathcal{P}_n$. Therefore,

$$\mathrm{Re}(W_\rho(\mathbf{u})) = \frac{1}{2^n}\mathrm{tr}\left[A_{\mathbf{u}}^\dagger \rho^{(0)}\right] = W_{\rho^{(0)}}(\mathbf{u}), \tag{A23}$$

$$\mathrm{Im}(W_\rho(\mathbf{u})) = \frac{1}{2^n}\mathrm{tr}\left[A_{\mathbf{u}}^\dagger \rho^{(1)}\right] = W_{\rho^{(1)}}(\mathbf{u}). \tag{A24}$$

■

An n-rebit Wigner representation $W_\rho^{(0)}$ was introduced by Delfosse *et al.* [44], which is defined as

$$W_\rho^{(0)}(\mathbf{u}) := \frac{1}{2^n}\mathrm{tr}\left[A_{\mathbf{u}}^{(0)\dagger}\rho\right], \tag{A25}$$

for all $\mathbf{u} \in \mathcal{P}_n$, where

$$A_{\mathbf{u}}^{(0)} := \frac{1}{2^n}\sum_{\mathbf{a}\in\mathcal{P}_n^0} (-1)^{[\mathbf{u},\mathbf{a}]} D_{\mathbf{a}} \tag{A26}$$

for the subset of phase space

$$\mathcal{P}_n^0 := \{\mathbf{u} \,:\, \mathbf{u}_x \cdot \mathbf{u}_z = 0\}. \tag{A27}$$

The definition of qubit displacement operators in Eq. (6) implies the relations

$$\mathbf{u}_x \cdot \mathbf{u}_z = 0 \Leftrightarrow D_{\mathbf{u}}^\dagger = D_{\mathbf{u}}^T = D_{\mathbf{u}}, \tag{A28}$$

$$\mathbf{u}_x \cdot \mathbf{u}_z = 1 \Leftrightarrow D_{\mathbf{u}}^\dagger = D_{\mathbf{u}}^T = -D_{\mathbf{u}}, \tag{A29}$$

We then see from Eq. (A28) that $\mathcal{P}_n^{(0)}$ is the subset of phase-space co-ordinates for *real symmetric* displacement operators. The difference between W_ρ and $W_\rho^{(0)}$ thus comes down to the fact that $A_{\mathbf{u}}$ sums over all displacement operators whereas $A_{\mathbf{u}}^{(0)}$ only sums over real symmetric ones, which implies $A_{\mathbf{u}}^{(0)}$ is itself real symmetric.

We further observe from Eq. (A29) that the complement of $\mathcal{P}_n^{(0)}$ in $\mathcal{P}_n$,

$$\mathcal{P}_n^{(1)} := \{\mathbf{u} \,:\, \mathbf{u}_x \cdot \mathbf{u}_z = 1\}, \tag{A30}$$

is the subset of phase-space co-ordinates for *real antisymmetric* displacement operators. Paralleling Eq. (A26), we then introduce the set of real antisymmetric phase-point operators

$$A_{\mathbf{u}}^{(1)} := \frac{1}{2^n}\sum_{\mathbf{a}\in\mathcal{P}_n^{(1)}} (-1)^{[\mathbf{u},\mathbf{a}]} D_{\mathbf{a}} \text{ for any } \mathbf{u} \in \mathcal{P}_n, \tag{A31}$$

By Eqs. (A26) and Eq. (A31), each $A_{\mathbf{u}}$ splits up as

$$A_{\mathbf{u}} = A_{\mathbf{u}}^{(0)} + A_{\mathbf{u}}^{(1)}. \tag{A32}$$

We can correspondingly split up the Wigner representation of ρ as

$$\begin{aligned} W_\rho(\mathbf{u}) &= \frac{1}{2^n}\mathrm{tr}[A_{\mathbf{u}}^\dagger\rho] = \frac{1}{2^n}\mathrm{tr}\left[\left(A_{\mathbf{u}}^{(0)} + A_{\mathbf{u}}^{(1)}\right)^\dagger \rho\right] \\ &= \frac{1}{2^n}\mathrm{tr}\left[A_{\mathbf{u}}^{(0)\dagger}\rho\right] + \frac{1}{2^n}\mathrm{tr}\left[A_{\mathbf{u}}^{(1)\dagger}\rho\right] \\ &=: W_\rho^{(0)}(\mathbf{u}) + W_\rho^{(1)}(\mathbf{u}), \end{aligned} \tag{A33}$$

where we define

$$W_\rho^{(1)}(\mathbf{x}) := \frac{1}{2^n}\mathrm{tr}\left[A_x^{(1)\dagger}\rho\right] = -\frac{1}{2^n}\mathrm{tr}\left[A_x^{(1)}\rho\right]. \tag{A34}$$

We can then prove the following.

Lemma 5.—Given any n-qubit state ρ,

$$Re(W_\rho(\mathbf{u})) = W_\rho^{(0)}(\mathbf{u}), \tag{A35}$$

$$iIm(W_\rho(\mathbf{u})) = W_\rho^{(1)}(\mathbf{u}). \tag{A36}$$

Proof.—Because $A_{\mathbf{u}}^{(0)}$ and $A_{\mathbf{u}}^{(1)}$ are, respectively, real symmetric and real antisymmetric, we have $A_{\mathbf{u}}^{(0)\dagger} = A_{\mathbf{u}}^{(0)T} = A_{\mathbf{u}}^{(0)}$ and $A_{\mathbf{u}}^{(1)\dagger} = A_{\mathbf{u}}^{(1)T} = -A_{\mathbf{u}}$. Thus for $k = 0, 1$, we obtain

$$\begin{aligned} \left[W_\rho^{(k)}(\mathbf{u})\right]^* &= \frac{1}{2^n}\mathrm{tr}\left[A_{\mathbf{u}}^{(k)\dagger}\rho\right]^* \\ &= \frac{1}{2^n}\mathrm{tr}\left[(A_{\mathbf{u}}^{(k)\dagger}\rho)^\dagger\right] \\ &= \frac{1}{2^n}\mathrm{tr}\left[(-1)^k(A_{\mathbf{u}}^{(k)\dagger}\rho)\right] \\ &= (-1)^k W_\rho^{(k)}(\mathbf{u}). \end{aligned} \tag{A37}$$

This implies $W_\rho^{(0)}(\mathbf{u})$ is the real component of $W_\rho(\mathbf{u})$ while $W_\rho^{(1)}(\mathbf{u})$ is its imaginary component. ■

By combining Lemmas 1 and 5, we arrive at

$$W_{\mathrm{Re}[\rho]}(\mathbf{u}) = W_{\mathbf{u}}^{(0)}(\rho). \tag{A38}$$

When ρ is an n-rebit state, $\mathrm{Re}(\rho) = \rho$, which implies $W_\rho(\mathbf{u}) = W_\rho^{(0)}(\mathbf{u})$.

3. Wigner representation of qubit channels

We recall from the main text that the Wigner representation of a channel $\mathcal{E} : \mathcal{B}(\mathcal{H}_2^n) \to \mathcal{B}(\mathcal{H}_2^m)$ is the linear map $W_\mathcal{E} : \mathcal{P}_n \to \mathcal{P}_m$ on phase space defined as

$$W_\mathcal{E}(\mathbf{v}|\mathbf{u}) := 2^{2n} W_{\mathcal{J}(\mathcal{E})}(\mathbf{u} \oplus \mathbf{v}), \tag{A39}$$

for all $\mathbf{v} \in \mathcal{P}_m, \mathbf{u} \in \mathcal{P}_n$, where the Choi state [51] of $\mathcal{E}$, $\mathcal{J}(\mathcal{E})$, is defined as $\mathcal{J}(\mathcal{E}) = (\mathcal{I} \otimes \mathcal{E}) \left|\phi_n^+\right\rangle\!\left\langle\phi_n^+\right|$ for the canonical maximally entangled state $\left|\phi_n^+\right\rangle$ on two sets of n qubits,

$$\left|\phi_n^+\right\rangle := \frac{1}{\sqrt{2n}} \left(\sum_{\mathbf{k} \in \{0,1\}^n} |\mathbf{k}\rangle \otimes |\mathbf{k}\rangle \right). \tag{A40}$$

One can straightforwardly verify that $\left|\phi_n^+\right\rangle$ is stabilized by $\langle Z_i Z_{n+i}, X_i X_{n+i} \rangle_{i=1,\ldots,n}$ and is therefore a CSS state.

The factorization property [Eq. (A1)] of phase-point operators implies that

$$W_\mathcal{E}(\mathbf{v}|\mathbf{u}) = \frac{2^n}{2^m} \mathrm{tr}\left[\left(A_\mathbf{u}^\dagger \otimes A_\mathbf{v}^\dagger\right) \mathcal{J}(\mathcal{E})\right]. \tag{A41}$$

Using the identity $\mathcal{E}(X) = 2^n \mathrm{tr}_{1,\ldots,n}\left[(X^T \otimes \mathbb{1}^{\otimes m})\mathcal{J}(\mathcal{E})\right]$ for transposition taken with respect to the computational basis, and recalling that $A_\mathbf{u}$ is real in the computational basis, we then conclude

$$W_\mathcal{E}(\mathbf{v}|\mathbf{u}) = \frac{1}{2^m} \mathrm{tr}[A_\mathbf{v}^\dagger \mathcal{E}(A_\mathbf{u}^*)] = \frac{1}{2^m} \mathrm{tr}[A_\mathbf{v}^\dagger \mathcal{E}(A_\mathbf{u})]. \tag{A42}$$

Therefore, if $\sigma = \mathcal{E}(\rho)$, then we obtain Eq. (17) from the main text, i.e.,

$$\begin{aligned}
W_\sigma(\mathbf{v}) &= \frac{1}{2^m} \mathrm{tr}\left[A_\mathbf{v}^\dagger \mathcal{E}(\rho)\right] \\
&= \frac{1}{2^m} \mathrm{tr}\left[\mathcal{E}\left(\sum_{\mathbf{u} \in \mathcal{P}_n} W_\rho(\mathbf{u}) A_\mathbf{u}\right) A_\mathbf{v}^\dagger\right] \\
&= \frac{1}{2^m} \sum_{\mathbf{u} \in \mathcal{P}_n} \mathrm{tr}\left[A_\mathbf{v}^\dagger \mathcal{E}\left(A_\mathbf{u}\right)\right] W_\rho(\mathbf{u}) \\
&= \sum_{\mathbf{u} \in \mathcal{P}_n} W_\mathcal{E}(\mathbf{v}|\mathbf{u}) W_\rho(\mathbf{u}).
\end{aligned} \tag{A43}$$

We thereby see that if $\mathcal{E}$ maps ρ to σ, then $W_\mathcal{E}$ is a matrix that maps W_ρ to W_σ, which justifies regarding $W_\mathcal{E}$ as the representation of $\mathcal{E}$ on phase space.

By property (A4) of the phase-point operators, we have that $\sum_{\mathbf{v} \in \mathcal{P}_m} A_\mathbf{v} = 2^m \mathbb{1}_m$. By applying this to the alternative formulation of $W_\mathcal{E}$ in Eq. (A42), we see that

$$\begin{aligned}
\sum_{\mathbf{v} \in \mathcal{P}_m} W_\mathcal{E}(\mathbf{v}|\mathbf{u}) &= \frac{1}{2^m} \sum_{\mathbf{v} \in \mathcal{P}_m} \mathrm{tr}\left[A_\mathbf{v}^\dagger \mathcal{E}(A_\mathbf{u})\right] \\
&= \frac{1}{2^m} \mathrm{tr}\left[\left(\sum_{\mathbf{v} \in \mathcal{P}_m} A_\mathbf{v}\right)^\dagger \mathcal{E}(A_\mathbf{u})\right] \\
&= \frac{1}{2^m} \mathrm{tr}[2^m \mathbb{1}_m \mathcal{E}(A_\mathbf{u})] \\
&= \mathrm{tr}[\mathcal{E}(A_\mathbf{u})].
\end{aligned} \tag{A44}$$

Then recalling that $\mathrm{tr}[A_\mathbf{u}] = 1$ [property (A3)], we obtain Eq. (21) from the main text, i.e.,

$$\sum_{\mathbf{v} \in \mathcal{P}_m} W_\mathcal{E}(\mathbf{v}|\mathbf{u}) = \mathrm{tr}[A_\mathbf{u}] = 1. \tag{A45}$$

This means every column of $W_\mathcal{E}$ sums up to 1.

Finally, we show that the representation we choose respects sequential and parallel composition of processes.

Let $\mathcal{E} : \mathcal{B}(\mathcal{H}_2^l) \to \mathcal{B}(\mathcal{H}_2^k)$ and $\mathcal{F} : \mathcal{B}(\mathcal{H}_2^n) \to \mathcal{B}(\mathcal{H}_2^m)$ be two multiqubit channels. Since $\{A_\mathbf{x}\}_{\mathbf{x} \in \mathcal{P}_m}$ are a complex orthogonal basis for $2^m \times 2^m$ complex matrices under the Hilbert-Schmidt inner product, we have that $\mathcal{F}(A_\mathbf{u}) = 1/2^m \sum_{\mathbf{x} \in \mathcal{P}_m} \mathrm{tr}[A_\mathbf{x}^\dagger \mathcal{F}(A_\mathbf{u})] A_\mathbf{x}$. Therefore, when $m = l$, we obtain

$$\begin{aligned}
W_{[\mathcal{E} \circ \mathcal{F}]}(\mathbf{v}|\mathbf{u}) &= \frac{1}{2^k} \mathrm{tr}[A_\mathbf{v}^\dagger \mathcal{E} \circ \mathcal{F}(A_\mathbf{u})] \\
&= \frac{1}{2^k} \mathrm{tr}\left[A_\mathbf{v}^\dagger \mathcal{E}\left(\frac{1}{2^m} \sum_{\mathbf{x} \in \mathcal{P}_m} \mathrm{tr}[A_\mathbf{x}^\dagger \mathcal{F}(A_\mathbf{u})] A_\mathbf{x}\right)\right] \\
&= \sum_{\mathbf{x} \in \mathcal{P}_m} \frac{1}{2^k} \mathrm{tr}[A_\mathbf{v}^\dagger \mathcal{E}(A_\mathbf{x})] \frac{1}{2^m} \mathrm{tr}[A_\mathbf{x}^\dagger \mathcal{F}(A_\mathbf{u})] \\
&= \sum_{\mathbf{x} \in \mathcal{P}_m} W_\mathcal{E}(\mathbf{v}|\mathbf{x}) W_\mathcal{F}(\mathbf{x}|\mathbf{u}),
\end{aligned} \tag{A46}$$

or in matrix notation,

$$W_{\mathcal{E} \circ \mathcal{F}} = W_\mathcal{E} W_\mathcal{F}. \tag{A47}$$

Furthermore, due to the factorization property of the phase-point operators, we have that

$$\begin{aligned}
W_{\mathcal{E} \otimes \mathcal{F}}(\mathbf{x} \oplus \mathbf{y}|\mathbf{u} \oplus \mathbf{v}) &= \frac{1}{2^{(m+k)}} \mathrm{tr}[A_\mathbf{x}^\dagger \otimes A_\mathbf{y}^\dagger \mathcal{E} \otimes \mathcal{F}(A_\mathbf{u} \otimes A_\mathbf{v})] \\
&= \frac{1}{2^k} \mathrm{tr}[A_\mathbf{x}^\dagger \mathcal{E}(A_\mathbf{u})] \frac{1}{2^m} \mathrm{tr}[A_\mathbf{y}^\dagger \mathcal{F}(A_\mathbf{v})] \\
&= W_\mathcal{E}(\mathbf{x}|\mathbf{u}) W_\mathcal{F}(\mathbf{y}|\mathbf{v}),
\end{aligned} \tag{A48}$$

or in matrix notation

$$W_{\mathcal{E}\otimes\mathcal{F}} = W_{\mathcal{E}} \otimes W_{\mathcal{F}}. \tag{A49}$$

APPENDIX B: COMPLETELY CSS-PRESERVING OPERATIONS

1. Completely CSS-preserving unitaries

The group of CSS-preserving unitaries on n qubits [44] can be generated as

$$\mathcal{G}_+(n) := \langle H^{\otimes n}, \textsc{cnot}(i,j), X_i, Z_i\rangle_{i,j=1,\ldots,n, i\neq j}\,. \tag{B1}$$

We also find it useful to note the following conjugation relations of the collective Hadamard gate:

$$H^{\otimes n}\,\textsc{cnot}(i,j) = \textsc{cnot}(j,i) H^{\otimes n} \tag{B2}$$

$$H^{\otimes n} X(\mathbf{a}) = Z(\mathbf{a}) H^{\otimes n}, \tag{B3}$$

which, respectively, hold for all $i,j \in \{1,\ldots n\}$ where $i \neq j$ and n-bit strings $\mathbf{a} \in \{0,1\}^n$.

Lemma 6.—The group of *completely CSS-preserving unitaries* on n qubits is

$$\mathcal{G}(n) := \langle \textsc{cnot}(i,j), Z_i, X_i\rangle_{i,j=1,\ldots,n, i\neq j}\,. \tag{B4}$$

Proof.—Let U_+ be any CSS-preserving unitary on n qubits. We first observe that U_+ is either in $\mathcal{G}(n)$ or is a unitary from $\mathcal{G}(n)$ followed by the collective Hadamard gate on n qubits, i.e., $U_+ = \left[H^{\otimes n}\right]^b U$ for some binary digit $b \in \{0,1\}$ and unitary $U \in \mathcal{G}(n)$. This follows from the conjugation relations Eqs. (B2) and Eq. (B3) alongside the fact that $H^{\otimes n}$ is self-inverse.

If U_+ is in $\mathcal{G}(n)$, then because $\mathcal{G}(n)$ is a subset of $\mathcal{G}_+(n')$ for all $n' \geq n$, U_+ must be *completely* CSS preserving. If U_+ is not in $\mathcal{G}(n)$, then we must have $U_+ = H^{\otimes n} U$ for some $U \in \mathcal{G}(n)$, which implies U_+ cannot be completely CSS preserving since $H^{\otimes n}$ is not completely CSS preserving. Therefore, U_+ is completely CSS preserving if and only if it is in $\mathcal{G}(n)$. ■

2. Completely CSS-preserving measurements

Throughout the rest of this Appendix, we extend, wherever necessary, the notion of being completely CSS preserving to trace-decreasing operations—i.e., a trace-decreasing operation $\mathcal{E}$ from n to n' qubits is completely CSS preserving if, given any CSS state ρ on $m+n$ qubits, we have that $(\mathcal{I}_m \otimes \mathcal{E})(\rho)$ is always a (possibly subnormalized) CSS state on $(m+n')$ qubits.

The projective measurement of any n-qubit Pauli observable S is carried out using projectors $P(\pm S) := 1/2\,(\mathbb{1}_n \pm S)$ corresponding to the ±1 outcomes. Postselection for the ±1 outcome is then carried out by the operation $\mathcal{P}(\pm S) := P(\pm S)(\cdot)P(\pm S)$.

Lemma 7.—Postselecting the ±1 outcome in the projective measurement of a CSS observable is completely CSS preserving.

Proof.—Let S be a CSS observable on n qubits and $|\psi\rangle$ be a CSS state on $m+n$ qubits for any $m \geq 0$. Then let $S_1,\ldots,S_{m+n}$ be a set of $m+n$ CSS observables that generate the stabilizer group $\mathcal{S}(|\psi\rangle)$ of $|\psi\rangle$.

Post-selecting the ±1 outcome in a projective measurement of S on the last n qubits of $|\psi\rangle$ yields the possibly subnormalized output

$$\begin{aligned}|\phi_\pm\rangle &:= \left[\mathbb{1}^{\otimes m} \otimes P(\pm S)\right]|\psi\rangle \\ &= \left[\frac{1}{2}\left(\mathbb{1}^{\otimes m+n} \pm \mathbb{1}_m \otimes S\right)\right]|\psi\rangle \\ &= P(\pm S')\,|\psi\rangle\,,\end{aligned} \tag{B5}$$

where we define the CSS observable $S' := \mathbb{1}^{\otimes m} \otimes S$. There are now two possibilities:

(a) that S' commutes with every generator of $\mathcal{S}(|\psi\rangle)$, so S' or $-S'$ must stabilize $|\psi\rangle$. Therefore, either $|\phi_+\rangle = |\psi\rangle$ and $|\phi_-\rangle = 0$ or vice versa, so $|\phi_\pm\rangle$ are possibly subnormalized CSS states;

(b) that, without loss of generality, S' does *not* commute with just one CSS observable S_1 that generates $\mathcal{S}(|\psi\rangle)$. This follows from the fact that, in any set of $m+n$ CSS observables that generate $\mathcal{S}(|\psi\rangle)$, those that do *not* commute with S' must *all* be X or Z type, so by picking one such generator and multiplying all others by it, we obtain another set of $m+n$ CSS observables generating $\mathcal{S}(|\psi\rangle)$ in which only one generator does *not* commute with S'. Then the states $|\phi_\pm\rangle$ have norm $1/\sqrt{2}$ and are stabilized, respectively, by $\langle \pm S, S_2, \ldots, S_{m+n}\rangle$, so $|\phi_\pm\rangle$ are subnormalized CSS states.

Therefore, given any pure CSS state $|\psi\rangle$ on $m+n$ qubits, postselecting the ±1 outcome in the projective measurement of a CSS observable on the last n qubits of $|\psi\rangle$ always produces a (possibly subnormalized) CSS state. Since every CSS state is a statistical mixture of pure CSS states, we arrive at the lemma result. ■

3. CSS circuits

In this section, we show that the subset of stabilizer operations covered by Lemma 2, which we referred to as *CSS circuits*, are completely CSS preserving.

To reiterate, a *CSS circuit* is any sequence of the following four *primitive CSS channels*:

(1) introducing a CSS state on any number of qubits,
(2) performing a completely CSS-preserving unitary,

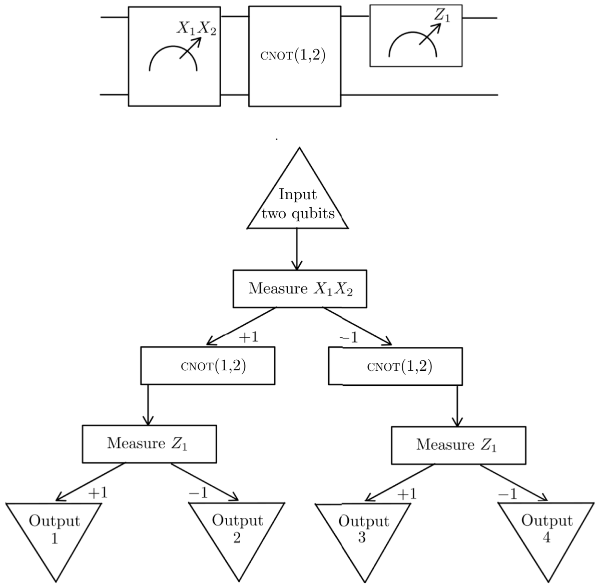

FIG. 7. The sequence of primitive CSS channels on the top can be executed as the binary tree on the bottom.

(3) projective measurement of any CSS observable, with the possibility of performing different sequences of primitive CSS channels depending on outcome,
(4) discarding any number of qubits,

as well as statistical mixtures of such sequences. We emphasize that all CSS circuits are trace preserving.

Any sequence of primitive CSS channels can be executed as a binary tree where the root node represents inputting qubits, the leaf nodes represent outputting qubits, and internal nodes represent primitive CSS channels. An illustrative example is provided in Fig. 7, from which we see that a sequence of primitive CSS channels can produce outputs distinguished by the sequences of measurement outcomes leading up to them.

Different numbers of ancillary qubits may be introduced on different branches of a tree representing a sequence of primitive CSS channels. However, one can arbitrarily increase the number of qubits introduced on any branch, without affecting what it does, by introducing the maximally mixed state on qubits that are immediately discarded just before the branch's output. Furthermore, different branches may have different lengths. However, one can arbitrarily lengthen any branch, without affecting what it does, by inserting identity channels just before the branch's output. Since introducing the maximally mixed state and the identity channel are both primitive CSS channels, we can, without loss of generality, consider only sequences of primitive CSS channels executed as binary trees where every branch has the same length and introduces the same number of ancillary qubits.

Lemma 8.—A CSS circuit $\mathcal{E}$ on n qubits is a statistical mixture of channels $\mathcal{E}_i$ representing sequences of primitive CSS channels, in which each sequence $\mathcal{E}_i$ is a sum of (possibly trace-decreasing) operations $\mathcal{E}_{i,j}$ generating its distinguishable outputs. Thus one can write

$$\begin{aligned}\mathcal{E}(\rho) &= \sum_i p_i \mathcal{E}_i, \text{ where } \mathcal{E}_i = \sum_j \mathcal{E}_{i,j},\\ \mathcal{E}_{i,j} &= \mathrm{tr}_{\mathcal{R}}\left[K_{i,j}(\rho\otimes\sigma_{i,j})K_{i,j}^\dagger\right], \text{ and}\\ K_{i,j} &= \prod_{l=1}^{N} P(S_{(i,j),l})U_{(i,j),l},\end{aligned} \tag{B6}$$

in which $\{p_i\}$ forms a probability distribution, $\sigma_{i,j}$ is a CSS state on m ancillary qubits, $\mathcal{R}$ is a subset of the $(n+m)$ input and ancillary qubits, $U_{(i,j),l}$ is a completely CSS-preserving unitary and $P(S_{(i,j),l})$ projects onto the $+1$ eigenspace of the CSS observable $S_{(i,j),l}$. Moreover, $P(S_{(i,j),1}),\ldots,P(S_{(i,j),N})$ gives the sequence of measurement outcomes that operationally distinguish the jth possible output of sequence $\mathcal{E}_i$.

Proof.—Without loss of generality, every branch of every sequence forming the mixture of $\mathcal{E}$ introduces the *same* number of ancillary qubits m and has *same* length N. One can show that the jth branch of the ith sequence must generate the channel $\mathcal{E}_{i,j}$ by induction over the steps of the branch. ■

Because each possible output from a sequence of primitive CSS channels results from a unique sequence of measurement outcomes, it is operationally meaningful to prepare a state conditioned upon obtaining the jth possible output in the ith sequence from the statistical mixture forming $\mathcal{E}$. We therefore have the following corollary.

Corollary 2.—We can record which operationally distinguishable output $\mathcal{E}_{i,j}$ from a CSS circuit $\mathcal{E}$ was obtained using a classical register,

$$\mathcal{E}'(\rho) := \sum_{i,j} p_i\, \mathcal{E}_{i,j}(\rho)\otimes \left|r_{i,j}\right\rangle\!\left\langle r_{i,j}\right|, \tag{B7}$$

where $\left|r_{i,j}\right\rangle$ is a computational basis state on multiple qubits. We note that, since $\left|r_{i,j}\right\rangle$ is a CSS state, $\mathcal{E}'$ can be carried out as a CSS circuit.

In the next subsection, we use this corollary to obtain the trace-preserving CSS-code projection studied throughout this paper.

We are finally in a position to prove Lemma 2, which is reproduced below. To this end, it is convenient to define the unique 0-qubit state 1 as CSS.

Lemma 2.—Any CSS circuit is completely CSS preserving.

Proof.—The decomposition of CSS circuits given in Lemma 8 implies that if (i) performing completely CSS-preserving unitaries, (ii) conditioning on the +1 outcome in the projective measurement of a CSS observable, (iii) introducing a CSS state, and (iv) discarding any number qubits are completely CSS preserving, then all CSS circuits are completely CSS preserving.

Now (i) is completely CSS preserving by definition, we proved that (ii) is completely CSS preserving in Lemma 7, and since the tensor product of two CSS states is always a CSS state, (iii) is completely CSS preserving.

Therefore, to prove that all CSS circuits are completely CSS preserving, we just have to prove that discarding any number of qubits is completely CSS preserving.

Consider discarding l qubits from n, where $n \geq l \geq 1$. Since we can freely relabel subsystems, we need to consider only discarding the *last* l qubits. Let $|\psi\rangle$ be a pure CSS state on $m+n$ qubits for any $m \geq 0$. Discarding the last l qubits of $|\psi\rangle$ then produces the state $\sigma := \mathcal{I}_m \otimes \mathrm{tr}_{n-l+1,\ldots,n}[|\psi\rangle\langle\psi|]$. Since tracing out is unaffected by first performing a computational basis measurement on the last l qubits, we have that

$$\sigma = \sum_{\mathbf{k}\in\{0,1\}^l} \mathrm{tr}_{m+n-l+1,\ldots,m+n}\Big[\mathbb{1}^{\otimes(m+n-l)} \otimes |\mathbf{k}\rangle\langle\mathbf{k}| \\ (|\psi\rangle\langle\psi|)\mathbb{1}^{\otimes(m+n-l)} \otimes |\mathbf{k}\rangle\langle\mathbf{k}|\Big]. \tag{B8}$$

We then observe that

$$\mathbb{1}^{\otimes(m+n-l)} \otimes |\mathbf{k}\rangle\langle\mathbf{k}| \\ = P((-1)^{k_1} Z_{m+n-l+1}) \circ \cdots \circ P((-1)^{k_l} Z_{m+n}), \tag{B9}$$

and so by Lemma 7, $\left(\mathbb{1}^{\otimes m+n-l} \otimes |\mathbf{k}\rangle\langle\mathbf{k}|\right)|\psi\rangle$ is a possibly subnormalized pure CSS state $\sqrt{p_\mathbf{k}}\,|\phi_\mathbf{k}\rangle \otimes |\mathbf{k}\rangle$, where $p_\mathbf{k}$ is the probability of getting the $|\mathbf{k}\rangle$ outcome in the computational basis measurement, and $|\phi_\mathbf{k}\rangle$ must be a (normalized) CSS state to keep the full state CSS. We thus obtain

$$\sigma = \sum_{\mathbf{k}\in\{0,1\}^l} \mathrm{tr}_{m+n-l+1,\ldots,m+n}\left(p_\mathbf{k}\,|\phi_\mathbf{k}\rangle\langle\phi_\mathbf{k}| \otimes |\mathbf{k}\rangle\langle\mathbf{k}|\right) \\ = \sum_{\mathbf{k}\in\{0,1\}^l} p_\mathbf{k}\,|\phi_\mathbf{k}\rangle\langle\phi_\mathbf{k}|, \tag{B10}$$

which is a CSS state on $(m+n-l)$ qubits. We conclude that, given any pure CSS state on $m+n$ qubits, discarding any $l \leq n$ of its final n qubits always produces a CSS state. This is true if and only if discarding any number of qubits is completely CSS preserving. ■

a. Omission of the collective Hadamard gate

The collective Hadamard gate promotes CSS circuits to a subset of stabilizer circuits where CSS states play the role of stabilizer states. One can reasonably ask why we exclude the collective Hadamard gate from the construction of CSS circuits. Our justification is that one can conjugate the collective Hadamard gate past any primitive CSS channel and leave another primitive CSS channel behind. This follows from the conjugation relations given by Eqs. (B2) and Eq. (B3) for completely CSS-preserving unitaries and projective measurements of CSS observables, from the cyclic property of the trace for discarding qubits, and from

$$H^{\otimes n} \otimes \mathbb{1}^{\otimes m}(\rho \otimes \sigma)H^{\otimes n} \otimes \mathbb{1}^{\otimes m} \\ = H^{\otimes(n+m)}(\rho \otimes (H^{\otimes m}\sigma H^{\otimes m}))H^{\otimes(n+m)} \tag{B11}$$

for introducing a CSS state on m ancillary qubits to an n-qubit system, where we note that $H^{\otimes m}\sigma H^{\otimes m}$ is also a CSS state because $H^{\otimes m}$ is CSS-preserving on m qubits. Therefore, circuits from this wider subset are operationally equivalent to CSS circuits followed by the collective Hadamard gate conditioned upon obtaining certain outputs, and are therefore not more powerful as magic distillation protocols.

4. CSS-code projections

An $[[n,k]]$ CSS code, where $n \geq 1$ and $n > k \geq 0$, is a vector space $\mathcal{C}$ stabilized by a subgroup $\mathcal{S}$ of n-qubit Pauli observables such that $-\mathbb{1}^{\otimes n} \notin \mathcal{S}$ and $\mathcal{S}$ can be generated from $n-k$ independent and commuting *CSS observables* $S_1, \ldots, S_{n-k}$, of which none is the identity.

Lemma 9.—Let $\mathcal{S} := \langle(-1)^{b_i}S_i\rangle_{i=1,\ldots,n-k}$ be the stabilizer group of an $[[n,k]]$ CSS code, where each S_i is a positive CSS observable and each b_i is a binary digit. Then there exists a completely CSS-preserving unitary U such that

$$U^\dagger[(-1)^{b_i}S_i]U = \begin{cases} Z_{k+i} & \text{if} S_i \text{is } Z, \\ X_{k+i} & \text{if } S_i \text{ is } X \text{ type.} \end{cases} \tag{B12}$$

Proof.—Let r be the number of Z-type generators for $\mathcal{S}$. We first construct the completely CSS-preserving unitary U for the case where all Z-type generators of $\mathcal{S}$ appear before X-type ones, i.e.,

$$S_i = \begin{cases} Z(\mathbf{u}_i) & \text{for } 1 \leq i \leq r \\ X(\mathbf{v}_i) & \text{for } r < i \leq n-k, \end{cases} \tag{B13}$$

where each $\mathbf{u}_i, \mathbf{v}_i$ is a nonzero n-dimensional binary vector.

Let us define $\mathcal{Z}$ as the set of Z-type generators for $\mathcal{S}$ *without their signs*, i.e., $\mathcal{Z} := \{Z(\mathbf{u}_1), \ldots, Z(\mathbf{u}_r)\}$. We now prove that there exists a sequence of CNOT operations that transforms $Z(\mathbf{u}_i)$ to Z_i for all $1 \leq i \leq r$.

We observe that $\mathcal{Z}$ is a subset of positive Z-type observables on n qubits, which form an n-dimensional vector space V over the field $\mathbb{F}_2$. Choosing the basis $\{Z_1, \ldots, Z_n\}$ for V, we can simply write $Z(\mathbf{u}_i)$ as $\mathbf{u}_i$, and we further have that $\{\mathbf{u}_1, \ldots, \mathbf{u}_r\}$ are linearly independent. Therefore, the matrix M_Z formed by taking members of $\mathcal{Z}$ as columns has rank r. Using Gauss-Jordan elimination, we can convert M_Z into its unique reduced row echelon form R_Z,

$$M_Z := \begin{bmatrix} | & \dots & | \\ \mathbf{u}_1 & \dots & \mathbf{u}_r \\ | & \dots & | \end{bmatrix} \mapsto R_Z := \begin{bmatrix} I_{r,r} \\ 0_{n-r,r} \end{bmatrix}, \tag{B14}$$

where $I_{r,r}$ is an $r \times r$ identity matrix while $0_{n-r,r}$ is a $(n-r) \times r$ null matrix.

On vector spaces over $\mathbb{F}_2$, Gauss-Jordan elimination consists of row swaps and additions. We now show how both can be done on any matrix M whose columns are elements of V in the basis $\{Z_1, \dots, Z_n\}$ using CNOT gates:

(1) *Swapping rows j and l.* This corresponds to swapping qubits j and l, which is carried out by performing $\text{CNOT}(l,j)\,\text{CNOT}(j,l)\,\text{CNOT}(a,b)$ on each positive Z-type observable forming a column in M.

(2) *Adding row j to row l.* The action of $\text{CNOT}(j,l)$ on Z_m is

$$\text{CNOT}(l,j) Z_m \,\text{CNOT}(l,j) = \begin{cases} Z_j Z_l & \text{for } m = j, \\ Z_m & \text{otherwise.} \end{cases} \tag{B15}$$

Therefore, given any vector $\mathbf{u}$ in V, we have that

$$\text{CNOT}(l,j) Z(\mathbf{u})\,\text{CNOT}(l,j) = Z(\mathbf{u} + u_j \mathbf{e}_l), \tag{B16}$$

where arithmetic is modulo 2 and $\mathbf{e}_l$ gives coordinates for the lth basis vector of V. In words, $\text{CNOT}(j,l)$ adds the jth component of $\mathbf{u}$ to the lth. Therefore, performing $\text{CNOT}(l,j)$ on each positive Z-type observable forming a column in M would add the jth row of M to the lth row.

We conclude that there exists a sequence of CNOT gates that, when performed on each element of $\mathcal{Z}$, accomplishes the Gauss-Jordan elimination in Eq. (B14). Denoting this sequence of CNOT gates by the unitary U_Z, we have that $U_Z(Z(\mathbf{u}_i))U_Z^\dagger = Z_i$ for all $1 \le i \le r$.

We next consider X-type generators for $\mathcal{S}$ *without their signs*. The action of $\text{CNOT}(j,l)$ on X_m is

$$\text{CNOT}(j,l) X_m \,\text{CNOT}(j,l) = \begin{cases} X_j X_l & \text{for } m = j, \\ X_m & \text{otherwise,} \end{cases} \tag{B17}$$

so U_Z transforms only positive X-type observables into other positive X-type observables. In other words, we can find nonzero n-bit strings $\{\mathbf{v}'_{r+1}, \dots, \mathbf{v}'_{n-k}\}$ such that $U_Z(X(\mathbf{v}_i))U_Z^\dagger = X(\mathbf{v}'_i)$ for all $r < i \le n-k$.

However, since the Z-type generators of $\mathcal{S}$ commute with the X-type generators, $X(\mathbf{v}'_i)$ must commute with $Z_1, \dots, Z_r$. Therefore, $X(\mathbf{v}'_i)$ must act trivially on qubits 1 through r, so the first r bits of $\mathbf{v}'_i$ must be 0. Therefore, $X(\mathbf{v}'_i) = \mathbb{1}^{\otimes r} \otimes X(\mathbf{v}''_i)$ for all $r < i \le n-k$, where $\mathbf{v}''_i$ are the last $(n-r)$ bits of $\mathbf{v}'_i$.

Everything we have done for $\mathcal{Z}$ can now be repeated for $\mathcal{X} := \{X(\mathbf{v}''_{r+1}), \dots, X(\mathbf{v}''_{n-k})\}$ on qubits $r+1$ through n. The only thing that needs to be checked is that row addition in any matrix whose columns are elements from the vector space of positive $(n-r)$-qubit X-type observables can be performed by CNOT gates. This can be confirmed using Eq. (B17), which implies

$$\text{CNOT}(j,l) X(\mathbf{v})\,\text{CNOT}(j,l) = X(\mathbf{v} + v_j \mathbf{e}_l), \tag{B18}$$

where arithmetic is modulo 2 so $\text{CNOT}(j,l)$ adds the jth component of $\mathbf{v}$ to the lth.

We conclude that there also exists a sequence U_X of CNOT gates such that $U_X(X(\mathbf{v}'_i))U_X^\dagger = X_i$ for all $r < i \le n-k$. Furthermore, U_X acts trivially on qubits 1 through r, which implies $U_X(Z_i)U_X^\dagger = Z_i$ for all $1 \le i \le r$. Defining $\mathcal{U}_X := U_X(\cdot)U_X^\dagger$ and $\mathcal{U}_Z := U_Z(\cdot)U_Z^\dagger$, we therefore have

$$\begin{aligned} &(\mathcal{U}_Z \circ \mathcal{U}_X)[(-1)^{b_i} S_i] \\ &= \begin{cases} (-1)^{b_i} X_i & \text{for } 1 \le i \le r \\ (-1)^{b_i} Z_i & \text{for } r < i \le n-k. \end{cases} \end{aligned} \tag{B19}$$

We now define unitaries U_C and U_{SWAP}, which, respectively, remove the signs from the generators of $\mathcal{S}$ and moves qubits 1 through $n-k$ to $k+1$ through n,

$$\begin{aligned} U_C &:= \left[\prod_{i=r+1}^{n-k} X_i^{b_i}\right]\left[\prod_{i=1}^{r} Z_i^{b_i}\right], \\ U_{\text{SWAP}} &:= \prod_{i=0}^{n-k-1} \text{SWAP}(n-i, n-k-i), \end{aligned} \tag{B20}$$

where $\text{SWAP}(i,j) := \text{CNOT}(i,j)\,\text{CNOT}(j,i)\,\text{CNOT}(i,j)$ swaps qubits i and j. Defining $\mathcal{U}_C := U_C(\cdot)U_C$ and $\mathcal{U}_{\text{SWAP}} := U_{\text{SWAP}}(\cdot)U_{\text{SWAP}}$, we obtain

$$\begin{aligned} &(\mathcal{U}_{\text{SWAP}} \circ \mathcal{U}_C \circ \mathcal{U}_Z \circ \mathcal{U}_X)[(-1)^{b_i} S_i] \\ &= \begin{cases} X_i & \text{for } 1 \le i \le r \\ Z_i & \text{for } r+1 \le i \le n-k. \end{cases} \end{aligned} \tag{B21}$$

Since $U^\dagger := U_C U_Z U_X$ is formed from CNOT-, single-qubit X and Z gates, it is a completely CSS-preserving unitary that accomplishes the lemma's claim for the ordering of S_i given by Eq. (B13).

By appropriate swaps among the last $n-k$ qubits, which we have seen is completely CSS preserving, we can construct a completely CSS-preserving unitary accomplishing the lemma's claim for any ordering of S_i. ■

As an immediate consequence of Lemma 9,

Corollary 3.—Every $[[n, k]]$ CSS code has a completely CSS-preserving encoding unitary.

Proof.—Let $\mathcal{C}$ be an $[[n, k]]$ CSS code whose stabilizer group is generated by CSS observables $S_1, \ldots, S_{n-k}$, and let $\mathbf{s}$ be a k-bit string (if $k = 0$, then $\mathbf{s}$ is the empty string). By Lemma 9, there exists a completely CSS-preserving unitary U such that

$$|\mathbf{s}_{\mathcal{C}}\rangle := U\left(|\mathbf{s}\rangle \bigotimes_{i=1}^{n-k} |\phi_i\rangle\right),$$

$$\text{where } |\phi_i\rangle := \begin{cases} |0\rangle & \text{if } S_i \text{ is } Z \text{ type} \\ |+\rangle & \text{if } S_i \text{ is } X \text{ type.} \end{cases} \tag{B22}$$

In words, U can encode any computational basis state $|\mathbf{s}\rangle$ on k physical qubits (if $k = 0$, then $|\mathbf{s}\rangle = 1$) as a *CSS logical basis state* $|\mathbf{s}_{\mathcal{C}}\rangle$ in $\mathcal{C}$. ■

With Corollary 3 in hand, we can show that every trace-preserving CSS-code projection can be executed as a CSS circuit, which leads to the following lemma from the main text.

Lemma 3.—Every trace-preserving CSS-code projection can be executed as a sequence of completely CSS preserving operations, and is therefore stochastically represented.

Proof.—Let $\mathcal{C}$ be an $[[n, k]]$ CSS code generated by CSS observables $S_1, \ldots, S_{n-k}$. We can then define a quantum channel $\mathcal{F}$ that carries out the following sequence of primitive CSS channels:

(1) Projectively measure $S_1, \ldots, S_{n-k}$, which is equivalent to the syndrome measurement of $\mathcal{C}$.
(2) If the no-error syndrome is obtained, decode onto the first k qubit using a completely CSS-preserving decoding unitary, which always exists in accordance with Corollary 3, and discard the final $n - k$ qubits.
(3) Otherwise, discard all qubits and prepare a k-qubit CSS state σ.

Since $\mathcal{F}$ is a CSS circuit by construction, it is completely CSS-preserving. By Corollary 2, we can convert $\mathcal{F}$ into a trace-preserving code projection $\mathcal{E}$ for $\mathcal{C}$ by distinguishing the output of the no-error syndrome from those of all other syndrome using a single-qubit classical register, i.e., by preparing an ancillary qubit in the state $|0\rangle$ if we obtain the no-error syndrome output, and in the state $|1\rangle$ otherwise. Corollary 2 confirms that $\mathcal{E}$ can also be executed as a CSS circuit and is therefore completely CSS-preserving. We conclude by Lemma 2 that $\mathcal{E}$ is stochastically represented. ■

APPENDIX C: ENTROPIC CONSTRAINTS ON COMPLETELY CSS-PRESERVING PROTOCOLS

In this section we prove Theorem 3, which gives entropic constraints on generic completely CSS-preserving protocols for qubits.

We first define the sets of positively-represented and real-represented quantum states in any Wigner representation W as

$$W^{+} := \{\rho : W_\rho(\mathbf{u}) > 0, W_\rho(\mathbf{u}) \in \mathbb{R}, \forall \mathbf{u} \in \mathcal{P}\}, \tag{C1}$$

$$W^{\mathbb{R}} := \{\rho : W_\rho(\mathbf{u}) \in \mathbb{R}, \forall \mathbf{u} \in \mathcal{P}\}, \tag{C2}$$

We then have the following Lemma, which is a generalisation of Theorem 11 of Ref. [27].

Lemma 10.—Let ρ and τ be states of a d-dimensional qudit such that $\rho \in W^{\mathbb{R}}$ and $\tau \in W^{+}$ in some generalised Gross's Wigner representation W. Furthermore, let $\mathcal{E} : \mathcal{B}(\mathcal{H}_d) \mapsto \mathcal{B}(\mathcal{H}_{d'})$ be a stochastically represented channel. Then the α-Rényi divergence $D_\alpha(\cdot||\cdot)$ is well-defined and satisfies the following properties for $\alpha \in \mathcal{A}$:

(1) $D_\alpha(W_\rho||W_\tau) \geq 0$.
(2) $D_\alpha(W_\rho||W_\tau) = 0$ if and only if $\rho = \tau$.
(3) $D_\alpha(W_{\rho^{\otimes n}}||W_{\tau^{\otimes n}}) = nD_\alpha(W_\rho||W_\tau)$ for all $n \in \mathbb{N}$.
(4) $D_\alpha(W_\rho||W_\tau) \geq D_\alpha(W_{\mathcal{E}(\rho)}||W_{\mathcal{E}(\tau)})$ for all stochastically represented $\mathcal{E}$ such that $\mathcal{E}(\tau) \in \mathcal{W}^{+}$.

Proof.—In general, W_ρ is a quasiprobability distribution, but for $\alpha \in \mathcal{A}$ we see that $W_\rho(\mathbf{u})^\alpha \geq 0$ for all $\mathbf{u} \in \mathcal{P}$. Therefore $D_\alpha(W_\rho||W_\tau)$ is always well-defined and real-valued. The proofs of 1–4 are then identical to the proof given for Theorem 11 in Ref.[27]. ■

Importantly for our purposes, this abstract but general result applies to input and output systems of any (even or odd) finite dimension. With this in hand, we can now give a proof of Theorem 3, which we restate for clarity:

Theorem 3.—Let ρ be a noisy rebit magic state and τ be a CSS state in the interior of $\mathcal{D}_{\text{CSS}}$. If there exists a completely CSS-preserving protocol $\mathcal{E}$ such that $\mathcal{E}(\rho^{\otimes n}) = \rho'$ and $\tau' := \mathcal{E}(\tau^{\otimes n})$ is also in the interior of $\mathcal{D}_{\text{CSS}}$, then

$$\Delta D_\alpha \geq 0 \tag{C3}$$

for all $\alpha \in \mathcal{A}$, where

$$\Delta D_\alpha := nD_\alpha(W_\rho||W_\tau) - D_\alpha(W_{\rho'}||W_{\tau'}). \tag{C4}$$

Proof.—Since τ and τ' are in the interior of $\mathcal{D}_{\text{css}}$, they are positively represented. Moreover, we have established in Theorem 1 that every trace-preserving and completely CSS-preserving operation is stochasically represented. The results of Lemma 10 thus apply, from which properties 3 and 4 combine to give the Lemma result. ■

APPENDIX D: ENTROPIC CONSTRAINTS ON CODE-PROJECTION PROTOCOLS

1. Structure of the necessary conditions

We recall from Sec. V that if there exists an n-to-k CSS-code projection that achieves the distillation $\rho^{\otimes n} \mapsto \rho'$ with acceptance probability p, then

$$\Delta D_\alpha \geq 0 \tag{D1}$$

for all $\alpha \in \mathcal{A}$, where

$$\Delta D_\alpha = nD_\alpha(W_\rho||W_{\frac{1}{2}}) - D_\alpha\left(W_{\rho_p} \,||\, W_{\tau_{n,k}}\right). \tag{D2}$$

We recall the following output states from the system and reference processes,

$$\rho_p := p\rho' \otimes |0\rangle\langle 0| + (1-p)\sigma \otimes |1\rangle\langle 1| \text{ and} \tag{D3}$$

$$\tau_{n,k} := 2^{k-n}\frac{\mathbb{1}^{\otimes k}}{2^k} \otimes |0\rangle\langle 0| + (1-2^{k-n})\sigma \otimes |1\rangle\langle 1|, \tag{D4}$$

where σ is the CSS state prepared after a failed run.

We now prove the following lemma, which enables us to simplify the expression of our constraint functions ΔD_α. To this end, we will find it useful to first introduce the general mean $Q_\alpha(\mathbf{w}||\mathbf{r})$ on a quasiprobabilty distribution $\mathbf{w} := (w_1, \dots w_N)^T$ and probability distribution $\mathbf{r} := (r_1, \dots r_N)^T$,

$$Q_\alpha(\mathbf{w}||\mathbf{r}) := 2^{(\alpha-1)D_\alpha(\mathbf{w}||\mathbf{r})} = \sum_{i=1}^{N} w_i^\alpha r_i^{1-\alpha}. \tag{D5}$$

Lemma 11.—Consider the following pairs of k-rebit states (ρ_0, ρ_1) and $(\tau_0, \tau_1) \in \mathrm{Int}(\mathcal{D}_{\mathrm{css}})$. Moreover let ψ_0 and ψ_1 be two distinct computational basis states on m rebits, and let $\{p_i\}_{i=0}^1$, $\{q_i\}_{i=0}^1$ be valid probability distributions. We then have the identity

$$Q_\alpha(W_{\sum_i p_i\rho_i\otimes\psi_i}||W_{\sum_i q_i\tau_i\otimes\psi_i}) = \sum_{i\in\{0,1\}} p_i^\alpha q_i^{1-\alpha} Q_\alpha(W_{\rho_i}||W_{\tau_i}), \tag{D6}$$

which in turn implies the inequality

$$Q_\alpha(W_{\sum_i p_i\rho_i\otimes\psi_i}||W_{\sum_i q_i\tau_i\otimes\psi_i}) \geq p_i^\alpha q_i^{1-\alpha} Q_\alpha(W_{\rho_i}||W_{\tau_i}), \tag{D7}$$

for each $i \in \{0,1\}$.

Proof.—Since ψ_0 and ψ_1 are orthogonal, by Eq. (A12) this implies

$$\mathrm{tr}[\psi_0\psi_1] = 2^m \sum_{\mathbf{u}\in\mathcal{P}_m} W_{\psi_0}(\mathbf{u})W_{\psi_1}(\mathbf{u}) = 0. \tag{D8}$$

As $\psi_i \in \mathcal{D}_{\mathrm{CSS}}$ for each $i \in \{0,1\}$, we must have $W_{\psi_i}(\mathbf{u}) \geq 0$. We thus conclude from Eq. (D8) that

$$V_0 := \mathrm{supp}(W_{\psi_0}) \subseteq \ker(W_{\psi_1}), \tag{D9}$$

$$V_1 := \mathrm{supp}(W_{\psi_1}) \subseteq \ker(W_{\psi_0}). \tag{D10}$$

With this in hand, we can explicitly evaluate:

$$\begin{aligned} &Q_\alpha(W_{\sum_i p_i\rho_i\otimes\psi_i}||W_{\sum_i q_i\tau_i\otimes\psi_i}) = \\ &\sum_{i\in\{0,1\}}\sum_{\mathbf{u}\in\mathcal{P}_k}\sum_{\mathbf{v}\in V_i}\left(p_iW_{\rho_i}(\mathbf{u})W_{\psi_i}(\mathbf{v})\right)^\alpha\left(q_iW_{\tau_i}(\mathbf{u})W_{\psi_i}(\mathbf{v})\right)^{1-\alpha} \\ &= \sum_{i\in\{0,1\}} p_i^\alpha q_i^{1-\alpha}\sum_{\mathbf{u}\in\mathcal{P}_k} W_{\rho_i}(\mathbf{u})^\alpha W_{\tau_i}(\mathbf{u})^{1-\alpha} \\ &= \sum_{i\in\{0,1\}} p_i^\alpha q_i^{1-\alpha} Q_\alpha(W_{\rho_i}||W_{\tau_i}), \end{aligned} \tag{D11}$$

where in the second equality we used the normalisation of our chosen representation W. The inequality in the Lemma statement then follows from the fact that both terms on the right hand side of Eq. (D6) must be non-negative for all $\alpha \in \mathcal{A}$. ■

With this property in hand, we obtain the following Lemma, which makes the non-trivial n-dependence in ΔD_α more explicit.

Lemma 12.—Let μ_k denote the maximally mixed state on k qubits. The function ΔD_α can then be expressed as

$$\begin{aligned}\Delta D_\alpha = n\left(1 - H_\alpha[W_\rho]\right) + k - \frac{1}{\alpha-1}\log \\ \times\left[p^\alpha Q_\alpha\left(W_{\rho'}||W_{\mu_k}\right) + (1-p)^\alpha\left(\frac{1}{2^{n-k}-1}\right)^{\alpha-1}\right],\end{aligned} \tag{D12}$$

Proof.—By Lemma 11, we have the expansion

$$\begin{aligned} Q_\alpha\left(W_{\rho_p}\,||\,W_{\tau_{n,k}}\right) &= p^\alpha\left(2^{k-n}\right)^{1-\alpha} Q_\alpha\left(W_{\rho'}\,||\,W_{\mu_k}\right) + (1-p)^\alpha\left(1-2^{k-n}\right)^{1-\alpha} Q_\alpha\left(W_\sigma\,||\,W_\sigma\right) \\ &= \left(2^{n-k}\right)^{\alpha-1}\left[p^\alpha Q_\alpha\left(W_{\rho'}||W_{\mu_k}\right) + (1-p)^\alpha\left(\frac{1}{2^{n-k}-1}\right)^{\alpha-1}\right], \end{aligned} \tag{D13}$$

where in the last equality we have used $Q_\alpha(\mathbf{p}||\mathbf{p}) = 1$ for all probability distributions $\mathbf{p}$. Therefore,

$$D_\alpha\left(W_{\rho_p}\,||\,W_{\tau_{n,k}}\right) = n - k + \frac{1}{\alpha-1}\log \left[p^\alpha Q_\alpha\left(W_{\rho'}||W_{\mu_k}\right) + (1-p)^\alpha\left(\frac{1}{2^{n-k}-1}\right)^{\alpha-1}\right]. \tag{D14}$$

Substituting Eq. (D14) and $D_\alpha\left(W_\rho\middle|\middle|W_{\mathbb{1}/2}\right) = 2 - H_\alpha[W_\rho]$ into Eq. (D2) gives the result as claimed. ■

2. Constraints are independent of the choice of CSS state σ prepared on failed code projection

By inspection, the form for ΔD_α given in Lemma 12 has no σ-dependence, which implies

Corollary 4.—The entropic constraints $\Delta D_\alpha \geq 0$ on n-to-k CSS-code-projection protocols are independent of which CSS state σ is prepared following failed runs.

This result also follows from resource-theoretic arguments. Consider a CSS circuit $\mathcal{E}$ on $k+1$ qubits that performs a Z-basis measurement on the last qubit and reprepares the first k qubits in a CSS state ω conditioned upon the -1 outcome. Thus one can write

$$\mathcal{E}_\omega(\cdot) := \mathcal{I}\otimes P_0(\cdot) + \omega\mathrm{tr}\otimes P_1(\cdot), \tag{D15}$$

where $P_k(\cdot) := |k\rangle\langle k|(\cdot)|k\rangle\langle k|$, from which one straightforwardly verifies that $\mathcal{E}_\omega(\rho_p)$ and $\mathcal{E}_\omega(\tau_{n,k})$ simply replaces σ in ρ_p and $\tau_{n,k}$ respectively by ω. Since $\mathcal{E}$ is stochastically represented, given any ω in the interior of $\mathcal{D}_{\mathrm{CSS}}$, we have by Property 4 in Lemma 10 that

$$\begin{aligned} D_\alpha\left(W_{\rho_p}||W_{\tau_{n,k}}\right) &\geq D_\alpha\left(W_{\mathcal{E}_\omega[\rho_p]}||W_{\mathcal{E}_\omega[\tau_{n,k}]}\right)\\ D_\alpha\left(W_{\rho_p}||W_{\tau_{n,k}}\right) &\geq D_\alpha\left(W_{\mathcal{E}_\sigma\circ\mathcal{E}_\omega[\rho_p]}||W_{\mathcal{E}_\sigma\circ\mathcal{E}_\omega[\tau_{n,k}]}\right)\\ &= D_\alpha\left(W_{\rho_p}||W_{\tau_{n,k}}\right). \end{aligned} \tag{D16}$$

We therefore conclude that

$$D_\alpha\left(\mathcal{E}_\omega(W_{\rho_p})||\mathcal{E}_\omega(W_{\tau_{n,k}})\right) = D_\alpha\left(W_{\rho_p}||W_{\tau_{n,k}}\right), \tag{D17}$$

so the entropic constraints on CSS-code projections are unaffected by varying σ in the interior of $\mathcal{D}_{\mathrm{css}}$.

3. Proof of Lemma 4

We now prove the following properties of the constraint function ΔD_α from the main text.

Lemma 4.—The following properties of the relative entropy difference ΔD_α hold for any noisy input rebit magic state ρ, output k-rebit magic state ρ', acceptance probability $p < 1$ and $\alpha \in \mathcal{A}$:

(i) ΔD_α is concave over the domain $n \in [k, \infty]$.

(ii) ΔD_α is negative in the limit where $n = k$:

$$\lim_{n\to k^+}\Delta D_\alpha < 0. \tag{D43}$$

.

(iii) If $H_\alpha[W_\rho] > 1$, then ΔD_α is also negative in the asymptotic limit

$$\lim_{n\to\infty}\Delta D_\alpha < 0. \tag{D44}$$

Proof.—Let us denote the maximally mixed state on k qubits by μ_k. We further simplify notation by defining the constants $c_1 := p^\alpha Q_\alpha(W_{\rho'}||W_{\mu_k})$ and $c_2 := (1-p)^\alpha$.

Proof of (i): Let us define the following function

$$g(n) := \left[c_1 + c_2\left(\frac{1}{2^{n-k}-1}\right)^{\alpha-1}\right]. \tag{D18}$$

This means that from Lemma 12 we can write

$$\Delta D_\alpha = n\left(1 - H_\alpha[W_\rho]\right) + k - \frac{1}{\alpha-1}\log g(n), \tag{D19}$$

and since the first term is linear we need only check the second derivative of the second term to establish that ΔD_α is concave. We have

$$\partial_n^2\Delta D_\alpha = -\frac{1}{\alpha-1}\partial_n^2\log g(n) = -\frac{\ln 2\, c_2 2^{k+n}\left(c_1\left(2^k + (\alpha-1)2^n\right)\left(2^{n-k}-1\right)^\alpha + c_2\left(2^n - 2^k\right)\right)}{\left(2^n-2^k\right)\left(c_1 2^k\left(2^{n-k}-1\right)^\alpha + c_2\left(2^n-2^k\right)\right)^2}. \tag{D20}$$

Since $c_1, c_2 \geq 0$ for all ρ' and p, the term in square brackets is non-negative for all $n > k, \alpha > 1, \rho'$ and p (strictly positive for $p < 1$), which implies $\partial_n^2\Delta D_\alpha$ is non-positive everywhere on our restricted domain. Therefore ΔD_α is concave, as claimed.

Proof of (ii). Recalling that $\alpha > 1$, we have from Lemma 12 that

$$\lim_{n\to k^+}\Delta D_\alpha = -kH_\alpha[W_{\rho'}]$$

$$
\begin{aligned}
&- \frac{1}{\alpha-1}\lim_{n\to k^+}\left\{\log\left[c_1+c_2\left(\frac{1}{2^{n-k}-1}\right)^{\alpha-1}\right]\right\}\\
&= -\infty < 0,
\end{aligned}
\tag{D21}
$$

so long as $c_2 > 1$, which is true if and only if $p < 1$.

Proof of (iii). We have

$$
\begin{aligned}
\lim_{n\to\infty}\Delta D_\alpha &= k - \frac{1}{\alpha-1}\log[c_1] + \lim_{n\to\infty}\left\{n\left(1-H_\alpha[W_\rho]\right)\right\}\\
&= H_\alpha[W_{p\rho'}] - k + \lim_{n\to\infty}\left\{n\left(1-H_\alpha[W_\rho]\right)\right\}\\
&= \begin{cases} -\infty, & H_\alpha[W_\rho] > 1,\\ +\infty, & H_\alpha[W_\rho] < 1,\\ H_\alpha[W_{p\rho'}]-k, & \text{otherwise.}\end{cases}
\end{aligned}
\tag{D22}
$$

Therefore, if $H_\alpha[W_\rho] > 1$ then $\lim_{n\to\infty}\Delta D_\alpha < 0$, as claimed.

This completes the proof. ■

4. Analytic bounds on code length in qudit code-projection protocols

In this section, we consider $[[n,k]]$ stabilizer-code projections for quantum systems of *arbitrary* Hilbert-space dimension d, which are stochastic under some generalized Gross's Wigner representation.

Following Eq. (36), we can define the following trace-preserving projection $\mathcal{E}$ on a $[[n,k]]$ stabilizer code $\mathcal{C}$ for d-dimensional quantum systems,

$$
\begin{aligned}
\mathcal{E}(\cdot) := &\operatorname{tr}_{k+1,\ldots,n}[\mathcal{U}\circ\mathcal{P}(\cdot)]\otimes|0\rangle\langle 0|\\
&+\operatorname{tr}[\overline{\mathcal{P}}(\cdot)]\sigma\otimes|1\rangle\langle 1|,
\end{aligned}
\tag{D23}
$$

where $\mathcal{U}$ and $\mathcal{P}$ are, respectively, the decoding channel and codespace projection for $\mathcal{C}$, $\overline{\mathcal{P}}$ is the projection onto the orthogonal complement of $\mathcal{C}$, and σ is positively represented under some generalized Gross's Wigner representation of our choice. This channel transforms n copies of an input noisy magic state ρ as

$$
\mathcal{E}\left[\rho^{\otimes n}\right] = p\rho'\otimes|0\rangle\langle 0| + (1-p)\sigma\otimes|1\rangle\langle 1| =: \rho_p,
\tag{D24}
$$

where we assume the output state ρ' following successful code projection is δ close to k copies of our pure target magic state ψ, as measured by the trace distance $\|\rho-\sigma\|_1$ where $\|X\|_1 := \operatorname{tr}|X| = \operatorname{tr}\sqrt{X^\dagger X}$ is the trace norm (also known as the Schatten-1 norm). Formally, we assume

$$
\|\rho' - \psi^{\otimes k}\|_1 \le \delta < \|\rho-\psi\|_1.
\tag{D25}
$$

We also define the Frobenius norm (also known as the Schatten-2 norm) as

$$
\|X\|_2 := \sqrt{\operatorname{tr}[X^\dagger X]}.
$$

We further define the ℓ_1- and ℓ_2-norms, respectively, of a vector $\mathbf{w}\in\mathbb{R}^d$ as

$$
\|\mathbf{w}\|_1 := \sum_{i=1}^{d}|w_i|,
\tag{D26}
$$

$$
\|\mathbf{w}\|_2 := \left[\sum_{i=1}^{d}|w_i|^2\right]^{\frac{1}{2}}.
\tag{D27}
$$

We now make use of the following result from the literature on real vector spaces (e.g., see Ref. [76]), which is a consequence of the Cauchy-Schwarz inequality.

Lemma 13.—For all $\mathbf{w}\in\mathbb{R}^d$

$$
\|\mathbf{w}\|_1 \le \sqrt{d}\|\mathbf{w}\|_2.
\tag{D28}
$$

This result enables us to show that vanishingly small variations in quantum states correspond to vanishingly small variations in their Wigner representations.

Lemma 14.—If $\|\rho-\sigma\|_1 \le \epsilon$ then for any generalised Gross's Wigner representation W, we have

$$
\|W_\rho - W_\sigma\|_1 \le \sqrt{d}\epsilon.
\tag{D29}
$$

Proof.—To simplify notation, we first define the state difference $\Delta := \rho-\sigma$ such that $W_\Delta = W_\rho - W_\sigma$. Since the Schatten-$p$ norms are nonincreasing with respect to p [76], we obtain

$$
\begin{aligned}
\|\rho-\sigma\|_1 \ge \|\rho-\sigma\|_2 &= \|\sum_{\mathbf{x}} W_\Delta(\mathbf{x})A_{\mathbf{x}}\|_2\\
&= \sqrt{\sum_{\mathbf{x},\mathbf{y}} W^*_\Delta(\mathbf{x})W_\Delta(\mathbf{y})\operatorname{tr}[A^\dagger_{\mathbf{x}}A_{\mathbf{y}}]}\\
&= \sqrt{\sum_{\mathbf{x},\mathbf{y}} W^*_\Delta(\mathbf{x})W_\Delta(\mathbf{y})d\delta_{\mathbf{x},\mathbf{y}}}\\
&= \sqrt{d}\|W_\Delta\|_2 \ge \frac{1}{\sqrt{d}}\|W_\Delta\|_1,
\end{aligned}
\tag{D30}
$$

where in the second inequality we employ Lemma 13. Therefore, $\|W_\rho - W_\sigma\|_1 \le \sqrt{d}\|\rho-\sigma\|_1$. ■

Lemma 15.—Let ρ and σ be two quantum states of a d-dimensional qudit such that $\|\rho-\sigma\|_1 \le \epsilon$. Then given any generalized Gross's Wigner representation W,

$$
|H_\alpha[W_\rho] - H_\alpha[W_\sigma]| \le \frac{\alpha}{\alpha-1}\log[1+\epsilon d^{\frac{5}{2}}].
\tag{D31}
$$

Proof.—Theorem 7 (2) of Ref. [77] applies to quasiprobability distributions and tells us that for two d^2-dimensional distributions $\mathbf{w},\mathbf{w}'$, we have the following continuity statement on the α-Rényi entropies:

$$|H_\alpha(\mathbf{w}) - H_\alpha(\mathbf{w}')| \leq \frac{\alpha}{\alpha-1}\log[1 + \|\mathbf{w}-\mathbf{w}'\|_1 d^2]. \tag{D32}$$

The proof of this just relies on the monotonicity of the p norms $\|\mathbf{w}\|_p := \left(\sum_{i=1}^{d^2}|w_i|^p\right)^{1/p}$, i.e., for $1 \leq \alpha < \beta \leq \infty$, $\|\mathbf{w}\|_\alpha \geq \|\mathbf{w}\|_\beta$, which also holds for quasidistributions. The lemma result then follows from Lemma 14. ■

We are now in a position to prove Theorem 5, which we restate here for clarity.

Theorem 5 (Qudit-code bounds).—Consider the distillation of k copies of a pure magic state ψ from a supply of the noisy magic state ρ, where ψ and ρ are d-dimensional qudit states that are real represented under a generalized Gross's Wigner representation W. Any stochastically represented distillation protocol that projects onto the codespace of an $[[n,k]]$ stabilizer code and can use n copies of ρ to distil out a k-qudit state ρ' with acceptance probability p and output error $\delta \geq \|\rho' - \psi^{\otimes k}\|_1$ must have a code length n such that

$$n \geq \frac{k\left[\log d - H_\alpha(W_\psi)\right] - \frac{\alpha}{1-\alpha}\log\left(\frac{p}{1+\delta d^{5/2}}\right)}{\left[\log d - H_\alpha(W_\rho)\right]}, \tag{D33}$$

for all $\alpha \in \mathcal{A}$ for which $H_\alpha(W_\rho) < \log d$, and

$$n \leq \frac{k\left[H_\alpha(W_\psi) - \log d\right] + \frac{\alpha}{1-\alpha}\log\left(\frac{p}{1+\delta d^{5/2}}\right)}{\left[H_\alpha(W_\rho) - \log d\right]}, \tag{D34}$$

for all $\alpha \in \mathcal{A}$ for which $H_\alpha(W_\rho) > \log d$.

Proof.—Let μ_k denote the maximally mixed state of k qudits with Hilbert-space dimension d. Following the same proof strategy as that of Lemma 11, we obtain

$$Q_\alpha(W_{\rho_p}||W_{\tau_{n,k}}) \geq \frac{p^\alpha}{d^{(n-k)(1-\alpha)}}Q_\alpha(W_{\rho'}||W_{\mu_k}). \tag{D35}$$

Since $\log(\cdot)/(\alpha-1)$ is a monotonically increasing function for $\alpha > 1$, it therefore follows that

$$\begin{aligned} D_\alpha(W_{\rho_p}||W_{\tau_{n,k}}) &\geq \\ \frac{1}{\alpha-1}&\log\left[\frac{p^\alpha}{d^{(n-k)(1-\alpha)}}Q_\alpha(W_{\rho'}||W_{\mu_k})\right] \\ &= D_\alpha(W_{\rho'}||W_{\mu_k}) + \frac{\alpha}{\alpha-1}\log p + (n-k)\log d \\ &= k\log d - H_\alpha(W_{\rho'}) + \frac{\alpha}{\alpha-1}\log p + n\log d, \end{aligned} \tag{D36}$$

where in the final equality we use the identity $D_\alpha(W_\rho||W_{\mathbb{1}_d/d}) = 2\log d - H_\alpha(W_\rho)$.

We can now make use of the continuity of the Rényi entropy as stated in Lemma 15 to further lower bound this divergence as

$$\begin{aligned} D_\alpha(W_{\rho_p}||W_{\tau_{n,k}}) \geq{}& k\left[\log d - H_\alpha(W_\psi)\right] \\ &- \frac{\alpha}{1-\alpha}\log\frac{p}{1+\delta d^{\frac{5}{2}}} + n\log d, \end{aligned} \tag{D37}$$

where the equality follows from the factorization of the α-Rényi entropy over subsystems. This gives rise to the following upper bound on the relative entropy difference $\Delta D_\alpha := nD_\alpha(W_\rho||W_{\mathbb{1}/d}) - D_\alpha(W_{\rho_p}||W_{\tau_{n,k}})$,

$$\begin{aligned} 0 \leq \Delta D_\alpha \leq{}& n[\log d - H_\alpha(W_\rho)] + k\left[H_\alpha(W_\psi) - \log d\right] \\ &+ \frac{\alpha}{1-\alpha}\log\frac{p}{1+\delta d^{\frac{5}{2}}}. \end{aligned} \tag{D38}$$

This gives a weaker but still necessary constraint on stochastic transformations accomplishing $\rho^{\otimes n} \mapsto \rho_p$ and $(\mathbb{1}/d)^{\otimes n} \mapsto \tau_{n,k}$, which we can rearrange as

$$\begin{aligned} n[H_\alpha(W_\rho) - \log d] \leq{}& k\left[H_\alpha(W_\psi) - \log d\right] \\ &+ \frac{\alpha}{1-\alpha}\log\frac{p}{1+\delta d^{\frac{5}{2}}}. \end{aligned} \tag{D39}$$

For $H_\alpha(W_\rho) < \log d$, we can rearrange Eq. (D37) to obtain the lower bound in Eq. (49). For $H_\alpha(W_\rho) > \log d$, we obtain the upper bound in Eq. (50). ■

APPENDIX E: DECOMPOSITION OF CSS MAGIC DISTILLATION INTO CODE PROJECTIONS

The goal of this Appendix is to prove Theorem 6, which generalizes Theorem 1 of Ref. [47] to an arbitrary number of output qubits in the CSS setting.

Theorem 6.—An *n-to-k CSS magic distillation protocol* is any CSS circuit that takes in $n \geq 2$ qubits and outputs onto the first $1 \leq k < n$ qubits. Every n-to-k CSS magic distillation protocol $\mathcal{E}$ can be decomposed as a sum of CSS-code projections, followed by preparing CSS states and completely CSS-preserving postprocessing. Thus one can write

$$\mathcal{E}(\rho) = \sum_j p_j\mathcal{E}_j,\ \mathcal{E}_j := \mathcal{U}_j \circ \left(\mathcal{K}_j(\rho) \otimes |\varphi_j\rangle\langle\varphi_j|\right), \tag{E1}$$

where p_j is a probability, $\mathcal{U}_j$ is a completely CSS-preserving unitary channel on k qubits, $\mathcal{K}_j$ is the codespace projection of an $[[n,k_j]]$ CSS code for some integer k_j in the range $0 \leq k_j \leq k$, and $|\varphi_j\rangle$ is a CSS state on $(k-k_j)$ qubits.

Proof.—To bring $\mathcal{E}$ into the desired form, we proceed in four steps (all auxiliary lemmas used will be presented after this proof):

(1) By Lemma 17, any n-to-k CSS magic distillation protocol $\mathcal{E}$ can be decomposed as a sum of n-qubit operations

$$\mathcal{E}(\rho) = \sum_i q_i \mathcal{E}_i(\rho), \tag{E2}$$

where q_i is a probability and

$$\mathcal{E}_i(\rho) := \mathrm{tr}_{k+1,\ldots,n+m}\left[K_i(\rho \otimes |\psi_i\rangle\langle\psi_i|)K_i^\dagger\right] \tag{E3}$$

for a CSS state $|\psi_i\rangle$ on m ancillary qubits and Kraus operator K_i. This Kraus operator has the form

$$K_i = U_i\left(\prod_{l=1}^{N} P(S_{i,l})\right), \tag{E4}$$

where U_i is a completely CSS-preserving unitary and $P(S_{i,l})$ is the projection onto the $+1$ eigenspace of a CSS observable $S_{i,l}$.

(2) By Lemma 18, each operation $\mathcal{E}_i$ can be decomposed into a sum

$$\mathcal{E}_i(\rho) = \sum_{\mathbf{s}\in\{0,1\}^{n+m-k}} \mathcal{E}_{i,\mathbf{s}}(\rho), \tag{E5}$$

where each operation $\mathcal{E}_{i,\mathbf{s}}$ first introduces m ancillary qubits in the CSS state $|\psi_i\rangle$, then postselects the $+1$ outcome in a sequence of projective measurements of CSS observables $S_{i,N},\ldots,S_{i,1}$, and then performs a CSS-code projection on the input and ancillary qubits. Thus one can write (note that $\mathcal{K}$ depends on i and $\mathbf{s}$)

$$\mathcal{E}_{i,\mathbf{s}}(\rho) = \mathcal{K}\circ\mathcal{P}(S_{i,1})\circ\cdots\circ\mathcal{P}(S_{i,N})[\rho\otimes|\psi_i\rangle\langle\psi_i|], \tag{E6}$$

in which $\mathcal{P}(S_{i,j})$ postselects the $+1$ outcome in a projective measurement of the CSS observable $S_{i,j}$, and $\mathcal{K}$ is the code projection of an $[[n+m,k]]$ CSS code.

(3) By repeated applications of Lemma 19, we find that $\mathcal{E}_{i,\mathbf{s}}$ performs a CSS-code projection on the input and ancillary qubits, followed by preparing a CSS state and completely CSS-preserving postprocessing. Thus one can write

$$\mathcal{E}_{i,\mathbf{s}}(\rho) = q'\,\mathcal{U}'\circ(\mathcal{K}'(\rho\otimes|\psi_i\rangle\langle\psi_i|)\otimes|\varphi'\rangle\langle\varphi'|), \tag{E7}$$

where q' is a probability, $\mathcal{U}'$ is a completely CSS-preserving unitary channel on k qubits, $|\varphi'\rangle$ is a CSS state on $k-k'$ qubits for some integer k' in the range $0\le k'\le k$, and $\mathcal{K}'$ is a code projection for an $[[n+m,k']]$ CSS code.

(4) By Lemma 21, each CSS-code projection $\mathcal{K}'$ on the input and ancillary qubits from Eq. (E7) can be reduced to a CSS-code projection on the input qubits *alone*, followed by preparing a CSS state and completely CSS-preserving postprocessing. Thus one can write

$$\mathcal{K}'(\rho\otimes|\psi_i\rangle\langle\psi_i|) = q''\,\mathcal{U}''\circ(\mathcal{K}''(\rho)\otimes|\varphi''\rangle\langle\varphi''|), \tag{E8}$$

where q'' is a probability, $\mathcal{U}''$ is a completely CSS-preserving unitary channel on k' qubits, $\mathcal{K}''$ is the code projection for an $[[n,k'']]$ CSS code for some integer k'' in the range $0\le k''\le k'$, and $|\varphi''\rangle$ is a CSS state on $k'-k''$ qubits. Substituting back then immediately yields the result.

■

1. Auxiliary lemmas

Before turning to the proofs of lemmas used in each step of the main proof, we first present a result that will be useful throughout.

Lemma 16.—Given any completely CSS-preserving unitary U and CSS observable S on n qubits, $S' := U^\dagger SU$ is another CSS observable of the same type as S. This further implies $P(\pm S)U = UP(\pm S')$.

Proof.—For convenience, let us label the n qubits as $1,\ldots,n$. Let $\mathbf{a}$ be an arbitrary n-bit string, and $\mathbf{e}_j$ be the n-bit string with 1 in its jth entry and 0 everywhere else. We then have the following conjugation relations:

$$Z_j[X(\mathbf{a})]Z_j = (-1)^{a_j}X(\mathbf{a}), \tag{E9}$$

$$Z_j[Z(\mathbf{a})]Z_j = Z(\mathbf{a}), \tag{E10}$$

$$X_j[X(\mathbf{a})]X_j = X(\mathbf{a}), \tag{E11}$$

$$X_j[Z(\mathbf{a})]X_j = X_j(-1)^{a_j}Z(\mathbf{a}), \tag{E12}$$

$$\mathrm{CNOT}(i,j)[X(\mathbf{a})]\,\mathrm{CNOT}(i,j) = X(\mathbf{a}+a_i\mathbf{e}_j), \tag{E13}$$

$$\mathrm{CNOT}(i,j)[Z(\mathbf{a})]\,\mathrm{CNOT}(i,j) = Z(\mathbf{a}+a_j\mathbf{e}_i), \tag{E14}$$

for any i,j from the range $1,\ldots,n$, where arithmetic is modulo 2. Since S is of the form $\pm X(\mathbf{a})$ or $\pm Z(\mathbf{a})$, and U is a product of $\mathrm{CNOT}(i,j)$, Z_j and X_j for i,j in the range $1,\ldots,n$, we immediately arrive at the lemma result. ■

It is furthermore useful to recall that, given an $[[n,k]]$ CSS code $\mathcal{C}$ whose stabilizer group is generated by CSS observables $S_1,\ldots,S_{n-k}$, Corollary 3 implies the code projection $\mathcal{K}$ of $\mathcal{C}$ can be represented using a completely

CSS-preserving encoding unitary U as

$$\mathcal{K}(\rho) := \mathrm{tr}_{k+1,\ldots,n}\left[U^\dagger P(\rho) P U\right],$$
$$\text{where } U^\dagger S_i U = \begin{cases} Z_{k+i} \text{ if } S_i \text{ is } Z \text{ type} \\ X_{k+i} \text{ if } S_i \text{ is } X \text{ type}\end{cases}$$
$$\text{and } P = \prod_{i=1}^{n-k} P(S_i). \tag{E15}$$

We remark that in the $k=0$ case, $\mathcal{K}$ simply projects onto a pure n-qubit CSS state and then discards it.

a. Step 1: Standard form for CSS magic distillation protocols

Lemma 17.—Any n-to-k CSS magic distillation protocol $\mathcal{E}$ can be decomposed as a sum of n-qubit operations

$$\mathcal{E}(\rho) = \sum_i p_i \mathcal{E}_i(\rho), \tag{E16}$$

where p_i is a probability and

$$\mathcal{E}_i(\rho) := \mathrm{tr}_{k+1,\ldots,n+m}\left[K_i(\rho \otimes |\psi_i\rangle\langle\psi_i|)K_i^\dagger\right], \tag{E17}$$

for a CSS state $|\psi_i\rangle$ on m ancillary qubits and Kraus operator K_i. Furthermore, K_i has the form

$$K_i = U_i\left(\prod_{l=1}^{N} P(S_{i,l})\right), \tag{E18}$$

where U_i is a completely CSS-preserving unitary and $P(S_{i,l})$ projects onto the $+1$ eigenspace of a CSS observable $S_{i,l}$.

Proof.—By Lemma 8, we can decompose $\mathcal{E}$ as follows:

$$\mathcal{E}(\rho) = \sum_i q_i \mathcal{E}_i(\rho), \tag{E19}$$

where q_i is a probability and

$$\mathcal{E}_i(\rho) = \mathrm{tr}_{k+1,\ldots,n+m}\left[K_i(\rho \otimes \sigma_i)K_i^\dagger\right] \tag{E20}$$

for a CSS state σ_i on m ancillary qubits and Kraus operator K_i. Furthermore, K_i has the form

$$K_i = \prod_{l=1}^{N} P(S_{i,l})U_{i,l}, \tag{E21}$$

where $U_{i,l}$ is a completely CSS-preserving unitary and $P(S_{i,l})$ projects onto the $+1$ eigenspace of the CSS observable $S_{i,l}$. We then conjugate every completely CSS-preserving unitary $U_{i,l}$ to the beginning of its Kraus operator K_i as shown in Lemma 16, where we compose them into a single completely CSS-preserving unitary U_i. Decomposing each CSS state σ_i on ancillary qubits into a statistical mixture of pure CSS states then yields the lemma result. ■

b. Step 2: Exposing the decoding

Lemma 18.—Consider a channel $\mathcal{E}$ on $n \geq 2$ qubits that performs a completely CSS-preserving unitary U and then discards the final $n-k$ qubits where $1 \leq k < n$,

$$\mathcal{E}(\rho) = \mathrm{tr}_{k+1,\ldots,n}\left[U^\dagger(\rho)U\right]. \tag{E22}$$

Then $\mathcal{E}$ can be decomposed into a sum over CSS-code projections $\mathcal{C}_{\mathbf{s}}$ indexed by the computational basis $\{\mathbf{s}\}$ on the discarded qubits,

$$\mathcal{E}(\rho) = \sum_{\mathbf{s}\in\{0,1\}^{n-k}} \mathcal{K}_{\mathbf{s}}(\rho). \tag{E23}$$

Proof.—The channel $\mathcal{E}$ is unchanged by performing a measurement in the computational basis $\{|\mathbf{s}\rangle\}$ of the final $n-k$ qubits before discarding them. Thus one can write

$$\mathcal{E}(\rho) = \sum_{\mathbf{s}\in\{0,1\}^{n-k}} \mathcal{K}_{\mathbf{s}}(\rho), \tag{E24}$$

where we define

$$\mathcal{K}_{\mathbf{s}}(\rho) := \mathrm{tr}_{k+1,\ldots,n}\left[\mathbb{1}^{\otimes k}\otimes|\mathbf{s}\rangle\langle\mathbf{s}|\,(U^\dagger(\rho)U)\mathbb{1}^{\otimes k}\otimes|\mathbf{s}\rangle\langle\mathbf{s}|\right]. \tag{E25}$$

Let $U_{\mathbf{s}}$ be a completely CSS-preserving unitary on the final $n-k$ qubits defined as $U_{\mathbf{s}}^\dagger := \bigotimes_{i=1}^{n-k}(X)^{s_i}$. By the cyclic property of the trace, we obtain

$$\mathcal{K}_{\mathbf{s}}(\rho) = \mathrm{tr}_{k+1,\ldots,n}\left[K(\rho)K^\dagger\right], \tag{E26}$$

for the Kraus operator

$$K := \mathbb{1}^{\otimes k}\otimes\left(U_{\mathbf{s}}^\dagger|\mathbf{s}\rangle\langle\mathbf{s}|\right)U^\dagger \tag{E27}$$
$$= \mathbb{1}^{\otimes k}\otimes|\mathbf{0}\rangle\langle\mathbf{0}|\left([\mathbb{1}^{\otimes k}\otimes U_{\mathbf{s}}^\dagger]U^\dagger\right). \tag{E28}$$

Conjugating $U' := U\left(\mathbb{1}^{\otimes k}\otimes U_{\mathbf{s}}\right)$ past $\mathbb{1}^{\otimes k}\otimes|\mathbf{0}\rangle\langle\mathbf{0}| = \prod_{i=1}^{n-k}P(Z_{k+i})$ in accordance with Lemma 16, we see by comparison with Eq. (E15) that $\mathcal{K}_{\mathbf{s}}$ is a code projection for an $[[n,k]]$ CSS code stabilized by $\langle U'Z_{k+1}U'^\dagger,\ldots,U'Z_nU'^\dagger\rangle$. ■

c. Step 3: Removing the projections

Lemma 19.—Let $\mathcal{E}$ be an operation on n qubits that projectively measures a CSS observable S, postselects the $+1$

outcome, and then carries out a code projection $\mathcal{K}$ of an $[[n,k]]$ CSS code $\mathcal{C}$. Thus one can write

$$\mathcal{E}(\rho) := \mathcal{K}\circ\mathcal{P}(S)(\rho),\ \mathcal{P}(S)(\cdot) := P(S)(\cdot)P(S). \tag{E29}$$

There are three possibilities for how $\mathcal{E}$ transforms ρ:

(1) that $\mathcal{E}(\rho) = 0$ for all ρ;

(2) that $\mathcal{E}$ is a code projection $\mathcal{K}'$ for another $[[n,k]]$ CSS code $\mathcal{C}'$, followed by completely CSS-preserving unitary postprocessing $\tilde{\mathcal{U}}$. Thus one can write

$$\mathcal{E}(\rho) = p\,\tilde{\mathcal{U}}\circ\mathcal{K}'(\rho), \tag{E30}$$

where $p > 0$ is a probability. One can further find logical bases $\left\{|\mathbf{s}_{\mathcal{C}'}\rangle\,|\mathbf{s}\in\{0,1\}^k\right\}$ and $\left\{|\mathbf{s}_{\mathcal{C}}\rangle\,|\mathbf{s}\in\{0,1\}^k\right\}$, respectively, generated by completely CSS-preserving encoding unitaries for the new code $\mathcal{C}'$ and the old code $\mathcal{C}$, such that

$$P(S)\,|\mathbf{s}_{\mathcal{C}}\rangle \propto |\mathbf{s}_{\mathcal{C}'}\rangle\,; \tag{E31}$$

(3) for $k\geq 1$, that $\mathcal{E}$ is a code projection $\mathcal{K}'$ for an $[[n,k-1]]$ CSS code $\mathcal{C}'$, followed by preparing a CSS state $|\varphi\rangle$ on a single qubit and completely CSS-preserving unitary postprocessing $\tilde{\mathcal{U}}$, i.e.,

$$\mathcal{E}(\rho) = \tilde{\mathcal{U}}\circ[\mathcal{K}'(\rho)\otimes|\varphi\rangle\langle\varphi|]. \tag{E32}$$

Furthermore, we can find logical bases $\left\{\left|\mathbf{s}'_{\mathcal{C}'}\right\rangle|\mathbf{s}\in\{0,1\}^{k-1}\right\}$ and $\left\{|\mathbf{s}_{\mathcal{C}}\rangle\,|\mathbf{s}\in\{0,1\}^k\right\}$, respectively, generated by completely CSS-preserving encoding unitaries for the new code $\mathcal{C}'$ and the old code $\mathcal{C}$, such that

$$P(S)\left|f\,(\mathbf{s}')_{\mathcal{C}}\right\rangle \propto \left|\mathbf{s}'_{\mathcal{C}'}\right\rangle, \tag{E33}$$

for some function $f\ :\{0,1\}^{k-1}\to\{0,1\}^k$.

Proof.—Let $S_{k+1},\dots,S_n$ be a set of $n-k$ CSS observables that generate the stabilizer group of $\mathcal{C}$. Beginning with the representation for the code projection $\mathcal{K}$ given by Eq. (E15), we can conjugate the encoding unitary past the codespace projector in accordance with Lemma 16, and arrive at the alternative representation

$$\mathcal{K}(\rho) = \mathrm{tr}_{k+1,\dots,n}\left[PU^\dagger(\rho)UP\right], \tag{E34}$$

where U is a completely CSS-preserving encoding unitary for $\mathcal{C}$ and P is now its no-error syndrome projector. Thus one can write

$$U^\dagger S_i U = C_i \text{ and } P := \prod_{i=k+1}^{n} P(C_i),\text{ where} \tag{E35}$$

$$C_i := \begin{cases} Z_i & \text{if } S_i \text{ is } Z \text{ type}\\ X_i & \text{if } S_i \text{ is } X \text{ type}\end{cases} \text{ for } i = k+1,\dots,n. \tag{E36}$$

Following Eq. (B22), U generates the following logical basis for $\mathcal{C}$:

$$\forall\mathbf{s}\in\{0,1\}^k : |\mathbf{s}_{\mathcal{C}}\rangle := U\left(|\mathbf{s}\rangle\bigotimes_{i=k+1}^{n}|\phi_i\rangle\right) \tag{E37}$$

$$\text{where } |\phi_i\rangle := \begin{cases} |0\rangle & \text{if } S_i \text{ is } Z \text{ type}\\ |+\rangle & \text{if } S_i \text{ is } X \text{ type.}\end{cases} \tag{E38}$$

By expressing P in Eq. (E34) as $P = \mathbb{1}^{\otimes k}\bigotimes_{i=1}^{n-k}|\phi_i\rangle\langle\phi_i|$ and expanding $\mathbb{1}^{\otimes k}$ in the computational basis, we arrive at another alternative form for $\mathcal{K}$,

$$\mathcal{K}(\rho) = K(\rho)K^\dagger,\text{ where } K^\dagger := \sum_{\mathbf{s}\in\{0,1\}^k}|\mathbf{s}_{\mathcal{C}}\rangle\langle\mathbf{s}_{\mathcal{C}}|\,\mathbf{s}. \tag{E39}$$

We now show how $\mathcal{E}$ can be manipulated into one of the forms stated by the lemma depending on the relationship between S and the CSS observables generating the stabilizer group for $\mathcal{C}$.

(i) ***S* does *not* commute with at least one generator $S_{k+1},\dots,S_n$.** We assume without loss of generality that S does not commute with S_{k+1}. We now show how $\mathcal{E}$ can be manipulated into the *second* form stated by the lemma.Using the form of $\mathcal{K}$ given in Eq. (E39), we can express $\mathcal{E}$ as

$$\mathcal{E}(\rho) = M(\rho)M^\dagger, \tag{E40}$$

where M is the Kraus operator

$$M^\dagger := \sum_{\mathbf{s}\in\{0,1\}^k}P(S)\,|\mathbf{s}_{\mathcal{C}}\rangle\langle\mathbf{s}_{\mathcal{C}}|\,\mathbf{s}. \tag{E41}$$

The stabilizer group $\mathcal{S}(|\mathbf{s}\rangle_{\mathcal{C}})$ for the logical basis state $|\mathbf{s}_{\mathcal{C}}\rangle$ in $\mathcal{C}$ can be generated from the set of n CSS observables $\{(-1)^{s_1}UZ_1U^\dagger,\dots,(-1)^{s_k}UZ_kU^\dagger, S_{k+1},\dots,S_n\}$. We can multiply all other members of this set that do *not* commute with S by S_{k+1} and, as all CSS observables that do not commute with S are of the same type, arrive at a set of n CSS observables $\{S_1, \dots, S_n\}$ such that

$$\mathcal{S}(|\mathbf{s}_{\mathcal{C}}\rangle) = \langle (-1)^{s_1} S_1, \dots, (-1)^{s_k} S_k, S_{k+1}, \dots, S_n \rangle. \tag{E42}$$

This implies

$$P(S)\,|\mathbf{s}_{\mathcal{C}}\rangle = \frac{1}{\sqrt{2}}\,|\psi_{\mathbf{s}}\rangle\,, \tag{E43}$$

where $|\psi_{\mathbf{s}}\rangle$ is a CSS state stabilized by

$$\mathcal{S}(|\psi_{\mathbf{s}}\rangle) = \langle (-1)^{s_1} S_1, \dots, (-1)^{s_k} S_k, S, S_{k+2} \dots, S_n \rangle. \tag{E44}$$

Applying Lemma 9 to $S_1, \dots, S_k, S, S_{k+2}, \dots, S_n$, we can find a completely CSS-preserving encoding unitary U' for the $[[n,k]]$ CSS code $\mathcal{C}'$ stabilized by $\langle S, S_{k+2}, \dots, S_n \rangle$ that generates a logical basis $\{|\mathbf{s}\rangle_{\mathcal{C}'}\}$ in $\mathcal{C}'$ such that $|\mathbf{s}_{\mathcal{C}'}\rangle$ shares the stabilizer group of $|\psi_{\mathbf{s}}\rangle$. Therefore, $|\mathbf{s}_{\mathcal{C}'}\rangle$ and $|\psi_{\mathbf{s}}\rangle$ differ only up to a phase, and one can write

$$P(S)\,|\mathbf{s}_{\mathcal{C}}\rangle = \frac{1}{\sqrt{2}} e^{-i\theta_{\mathbf{s}}}\,|\mathbf{s}_{\mathcal{C}'}\rangle\,. \tag{E45}$$

Substituting into Eq. (E40), we obtain

$$\mathcal{E}(\rho) = \frac{1}{2}\tilde{U}\left[\mathcal{K}'(\rho)\right]\tilde{U}^{\dagger}, \tag{E46}$$

where we define the following unitary on k qubits to adjust for the phase differences between $|\psi_{\mathbf{s}}\rangle$ and $|\mathbf{s}_{\mathcal{C}'}\rangle$,

$$\tilde{U} := \sum_{\mathbf{s}\in\{0,1\}^k} e^{i\theta_{\mathbf{s}}}\,|\mathbf{s}\rangle\langle\mathbf{s}|\,, \tag{E47}$$

as well as the operation

$$\mathcal{K}'(\cdot) := K'(\cdot)K'^{\dagger},\ K'^{\dagger} := \sum_{\mathbf{s}\in\{0,1\}^k} \left|\mathbf{s}'_{\mathcal{C}}\right\rangle\left\langle\mathbf{s}'_{\mathcal{C}}\right|\mathbf{s}. \tag{E48}$$

By comparison with Eq. (E39), we see that $\mathcal{K}'$ is the CSS-code projection of $\mathcal{C}'$. Equations Eq. (E45) and Eq. (E46) now match the second statement of the lemma provided that $\tilde{U}$ is completely CSS preserving, which we now proceed to show.We first express $\mathcal{E}$ using the form of $\mathcal{K}$ given by Eq. (E34) as

$$\mathcal{E}(\rho) = \mathrm{tr}_{k+1,\dots,n}\left[PU^{\dagger}P(S)\rho P(S)UP\right]. \tag{E49}$$

Let $|\psi\rangle$ be a CSS state on k qubits expressed as $|\psi\rangle = \sum_{\mathbf{s}\in\{0,1\}^k} c_{\mathbf{s}}\,|\mathbf{s}\rangle$ in the computational basis. Since U' is completely CSS preserving, it encodes $|\psi\rangle$ as another CSS state $|\psi_{\mathcal{C}'}\rangle = \sum_{\mathbf{s}} c_{\mathbf{s}}\,|\mathbf{s}_{\mathcal{C}'}\rangle$ in $\mathcal{C}'$. By inputting $|\psi_{\mathcal{C}'}\rangle\langle\psi_{\mathcal{C}'}|$ into the two forms of $\mathcal{E}$ given by Eqs. (E46) and Eq. (E49) and equating their outputs, we obtain

$$\frac{1}{2}\tilde{U}\,|\psi\rangle\langle\psi|\,\tilde{U}^{\dagger} = \mathrm{tr}_{k+1,\dots,n} \times \left[PU^{\dagger}P(S)(|\psi_{\mathcal{C}'}\rangle\langle\psi_{\mathcal{C}'}|)P(S)UP\right]. \tag{E50}$$

We see from the proof of Lemma 2 that the right-hand side carries out a completely CSS-preserving operation. Therefore, since $|\psi_{\mathcal{C}'}\rangle$ is CSS, $\tilde{U}\,|\psi\rangle$ must also be CSS. As this argument applies to *any* pure CSS state $|\psi\rangle$ on k qubits, we conclude that $\tilde{U}$ is CSS preserving.

From the proof of Lemma 6, we can express $\tilde{U}$ as $\tilde{U} = \left[H^{\otimes k}\right]^a V$, where V is a *completely* CSS-preserving unitary and a is a binary digit. While the CNOT gate and single-qubit X and Z gates map computational basis states onto computational basis states, the collective Hadamard gate does not. Since $\tilde{U}$ is diagonal in the computational basis, we conclude that $a = 0$, so $\tilde{U}$ is indeed completely CSS preserving.

(ii) ***S* commutes with $\boldsymbol{S_{k+1}, \dots, S_n}$ and is independent of them.** This is only possible when $k \geq 1$. In this case, we show that $\mathcal{E}$ can be manipulated into the *third* form stated in the lemma.By conjugating $P(S)$ past U in Eq. (E49) using Lemma 16, we can express $\mathcal{E}$ as

$$\mathcal{E}(\rho) = \mathrm{tr}_{k+1,\dots,n}\left[PP(S')U^{\dagger}(\rho)UP(S')P\right], \tag{E51}$$

where $S' := U^{\dagger}SU$ is another CSS observable.

Consider the term

$$P(S')P = P(S')P(C_{k+1})\dots P(C_n) \tag{E52}$$

from Eq. (E51). Since S' commutes with $C_{k+1}, \dots, C_n$, and C_i is a CSS observable on qubit i alone, S' must be trivial on any qubit i out of the last $n-k$ where C_i is a different type of CSS observable from S'.Given any two Pauli observables O_2 and O_2, we have that

$$P(O_1)P(O_2) = P(O_1O_2)P(O_2). \tag{E53}$$

We can therefore apply Eq. (E53) to S' and the C_i in Eq. (E52) to eliminate the part of $P(S')$ that acts on the last $n-k$ qubits, i.e., there exists a CSS observable S'' on the first k qubits *alone* such that

$$P(S')P = [P(S'') \otimes \mathbb{1}^{\otimes(n-k)}]P. \tag{E54}$$

Since S' must also be independent of $C_k + 1, \dots, C_n$, we conclude that S'' is not proportional to the identity.We now consider the cases where S'' is Z type and X type separately. To this end, it is first useful to define a completely CSS-preserving k-qubit unitary that moves the jth qubit in a block of k to the kth position,

$$\mathrm{MV}(j \to k) := \begin{cases} \prod_{i=j}^{k} \mathrm{SWAP}(i, i+1) & \text{if } j < k \\ \mathbb{1}^{\otimes k} & \text{otherwise,} \end{cases} \tag{E55}$$

where $\mathrm{SWAP}(i,j)$ swaps qubits i and j and is carried out by $\mathrm{CNOT}(i,j)\,\mathrm{CNOT}(j,i)\,\mathrm{CNOT}(i,j)$.

(a) ***S″* is *Z* type.** We therefore represent S'' as $S'' = (-1)^b Z(\mathbf{z})$, where b is a binary digit and $\mathbf{z}$ is an k-bit string of which at least one bit j is such that $z_j = 1$.Consider the following completely CSS-preserving unitary on k qubits

$$\tilde{U}^\dagger := \mathrm{MV}(j \to k) \circ X_j^b \circ U_C, \tag{E56}$$

where U_C is defined as

$$U_C := \begin{cases} \prod^{k}{}_{\substack{i=1 \\ \mathbf{z}_i=1, i\neq j}} \mathrm{CNOT}(i,j) & \text{if } \exists i \neq j : z_i = 1, \\ \mathbb{1}^{\otimes k} & \text{otherwise.} \end{cases} \tag{E57}$$

Because U_C acts as $U_C[Z(\mathbf{z})]U_C^\dagger = Z_j$, we have that

$$\tilde{U}^\dagger P(S'')\tilde{U} = P(Z_k), \tag{E58}$$

Using Eq. (E58) to substitute $P(S'')$ in Eq. (E54), we have by Eq. (E51) that

$$\mathcal{E}(\rho) = \tilde{U}\left(\mathrm{tr}_{k+1,\ldots,n}\left[P'U'^\dagger(\rho)U'P'\right]\right)\tilde{U}^\dagger, \tag{E59}$$

where we define

$$P' := P(Z_k)P \text{ and } U' := U\tilde{U} \otimes \mathbb{1}^{\otimes(n-k)}. \tag{E60}$$

Since kth qubit outputted by the part inside $\tilde{U}(\cdot)\tilde{U}^\dagger$ is always $|0\rangle$, we can simply discard the kth qubit as well and reprepare it. Thus one can write

$$\mathcal{E}(\rho) = \tilde{U}\left(\mathcal{K}'(\rho) \otimes |0\rangle\langle 0|\right)\tilde{U}^\dagger, \tag{E61}$$

where we define

$$\mathcal{K}'(\rho) := \mathrm{tr}_{k,\ldots,n}\left[P'U'^\dagger(\rho)U'P'\right]. \tag{E62}$$

From Eq. (E34), we see that $\mathcal{K}'$ is the code projection of an $[[n, k-1]]$ CSS code $\mathcal{C}'$ stabilized by $\langle U'Z_kU'^\dagger, U'C_{k+1}U'^\dagger, \ldots, U'C_nU'^\dagger\rangle$ with a completely CSS-preserving encoding unitary U'. By Eq. (E37), U' generates the following logical basis for $\mathcal{C}'$:

$$\forall \mathbf{s}' \in \{0,1\}^{k-1} : \left|\mathbf{s}'_{\mathcal{C}}\right\rangle := U'\,|\mathbf{s}\rangle\,|0\rangle\,|\Phi\rangle \tag{E63}$$

$$\text{where } |\Phi\rangle := \bigotimes_{i=1}^{n-k} |\phi_i\rangle. \tag{E64}$$

We can then relate the logical basis state $\left|\mathbf{s}'_{\mathcal{C}'}\right\rangle$ to a logical basis state in the old code $\mathcal{C}$ as

$$\begin{aligned} \left|\mathbf{s}'_{\mathcal{C}'}\right\rangle &= \left[U\tilde{U} \otimes \mathbb{1}^{\otimes(n-k)}\right] P(Z_k)P\left|\mathbf{s}'\right\rangle|0\rangle\,|\Phi\rangle \\ &= UP\left(S''\right) \otimes \mathbb{1}^{\otimes(n-k)}P\left[\left(\tilde{U}\left|\mathbf{s}'\right\rangle|0\rangle\right)|\Phi\rangle\right] \\ &= UP(S')P\left|f(\mathbf{s}')\right\rangle|\Phi\rangle \\ &= P(S)U\left|f(\mathbf{s}')\right\rangle|\Phi\rangle \\ &= P(S)\left|f(\mathbf{s}')_{\mathcal{C}}\right\rangle, \end{aligned} \tag{E65}$$

where we define $\left|f(\mathbf{s}')\right\rangle := \tilde{U}\left|\mathbf{s}'\right\rangle|0\rangle$, used Eq. (E58) to obtain the second equality and Eq. (E54) to obtain the third. Explicitly, f is defined by

$$\begin{gathered} f(\mathbf{s}') = (s'_1, \ldots, s'_{j-1}, s, s'_j, \ldots, s'_{k-1}), \text{where} \\ s := \left(\sum_{i=1}^{j-1} s'_i z_i\right) + b + \left(\sum_{i=j+1}^{k} s'_{i-1} z_i\right), \end{gathered} \tag{E66}$$

and arithmetic is modulo 2.

(b) ***S″* is *X* type.** We therefore represent S'' as $S'' = (-1)^b X(\mathbf{x})$, where b is a binary digit and $\mathbf{x}$ is an k-bit string of which at least one bit j is such that $x_j = 1$.Consider the following completely CSS-preserving unitary on k qubits,

$$\tilde{U}^\dagger := \mathrm{MV}_{j\to k} \circ Z_j^b \circ U_C, \tag{E67}$$

in which U_C is defined as

$$U_C := \begin{cases} \prod^{k}{}_{\substack{i=1 \\ \mathbf{x}_i=1, i\neq j}} \mathrm{CNOT}(j,i) & \text{if } \exists i \neq j : x_i = 1, \\ \mathbb{1}^{\otimes k} & \text{otherwise.} \end{cases} \tag{E68}$$

Because $U_C[X(\mathbf{x})]U_C^\dagger = X_j$, we have that

$$\tilde{U}^\dagger P(S'')\tilde{U} = P(X_k). \tag{E69}$$

Using Eq. (E69) to substitute $P(S'')$ in Eq. (E54), we have by Eq. (E51) that

$$\mathcal{E}(\rho) = \tilde{U}\left(\mathrm{tr}_{k+1,\ldots,n}\left[P'U'^{\dagger}(\rho)U'P'\right]\right)\tilde{U}^{\dagger}, \tag{E70}$$

where we define

$$P' := P(X_k)P,\ U' := U\tilde{U}\otimes \mathbb{1}^{\otimes(n-k)}. \tag{E71}$$

Since kth qubit outputted by the part inside $\tilde{U}(\cdot)\tilde{U}^{\dagger}$ is always $|+\rangle$, we can simply discard the kth qubit as well and reprepare it. Thus one can write

$$\mathcal{E}(\rho) = \tilde{U}\left(\mathcal{K}'(\rho)\otimes|+\rangle\langle+|\right)\tilde{U}^{\dagger}, \tag{E72}$$

where we define

$$\mathcal{K}'(\rho) := \mathrm{tr}_{k,\ldots,n}\left[P'U'^{\dagger}(\rho)U'P'\right]. \tag{E73}$$

From Eq. (E15), we see that $\mathcal{K}'$ is a code projection for an $[[n,k-1]]$ CSS code $\mathcal{C}'$ stabilized by $\langle U'X_kU'^{\dagger}, U'C_{k+1}U'^{\dagger},\ldots,U'C_nU'^{\dagger}\rangle$ with a completely CSS-preserving encoding unitary U'. By Eq. (E37), U' generates the following logical basis for $\mathcal{C}'$:

$$\forall \mathbf{s}'\in\{0,1\}^{k-1} : \left|\mathbf{s}'_{\mathcal{C}}\right\rangle := U'\,|\mathbf{s}\rangle\,|+\rangle\,|\Phi\rangle \tag{E74}$$

$$\text{where } |\Phi\rangle := \bigotimes_{i=1}^{n-k}|\phi_i\rangle\,. \tag{E75}$$

We can then relate the logical basis state $\left|\mathbf{s}'_{\mathcal{C}'}\right\rangle$ in the new code $\mathcal{C}'$ to a logical basis state in the old code $\mathcal{C}$ as

$$\begin{aligned}\left|\mathbf{s}'_{\mathcal{C}'}\right\rangle &= \left[U\tilde{U}\otimes\mathbb{1}^{(n-k)}\right]P(X_k)P\left|\mathbf{s}'\right\rangle|+\rangle\,|\Phi\rangle\\ &= \sqrt{2}\left[U\tilde{U}\otimes\mathbb{1}^{\otimes(n-k)}\right]P(X_k)P\left|\mathbf{s}'\right\rangle|0\rangle\,|\Phi\rangle\\ &= \sqrt{2}\,UP(S'')\otimes\mathbb{1}^{\otimes(n-k)}P\left(\tilde{U}\left|\mathbf{s}'\right\rangle|0\rangle\right)|\Phi\rangle\\ &= \sqrt{2}\,UP(S')P\left|f(\mathbf{s}')\right\rangle|\Phi\rangle\\ &= \sqrt{2}\,P(S)U\left|f(\mathbf{s}')\right\rangle|\Phi\rangle\\ &= \sqrt{2}\,P(S)\left|f(\mathbf{s}')_{\mathcal{C}}\right\rangle,\end{aligned} \tag{E76}$$

where we define $\left|f(\mathbf{s}')\right\rangle := \tilde{U}\left|\mathbf{s}'\right\rangle|0\rangle$, used Eq. (E69) to obtain the second equality and Eq. (E54) to obtain the third. Explicitly, $f(\mathbf{s}') = (s'_1,\ldots,s'_{j-1},0,s'_j,\ldots,s'_{k-1})$.

(iii) ***S* is not independent of $S_{k+1},\ldots,S_n$.** This implies either $-S$ or S stabilizes $\mathcal{C}$. In the former case, we see from Eq. (E40) that $\mathcal{E}(\rho)=0$ for all ρ. In the latter case, we have $P(S)\,|\mathbf{s}_{\mathcal{C}}\rangle = |\mathbf{s}_{\mathcal{C}}\rangle$, which implies $\mathcal{E}=\mathcal{K}$ by Eq. (E40). Together these equations match the lemma's second form.

■

d. Step 4: Removing ancillary qubits

Lemma 20.—Let $\mathcal{C}$ be an $[[n+m,k]]$ CSS codes where $n\geq 1, n>k$ and $m>0$, and let $\{|\mathbf{s}_{\mathcal{C}}\rangle\}$ be a logical basis for $\mathcal{C}$ generated by a completely CSS-preserving encoding unitary U such that each logical basis state $|\mathbf{s}_{\mathcal{C}}\rangle$ factorizes over n and m qubits, i.e.,

$$|\mathbf{s}_{\mathcal{C}}\rangle = |\psi_{\mathbf{s}}\rangle\otimes|\psi\rangle\,, \tag{E77}$$

where $|\psi\rangle$ is a CSS state on m qubits. We then have that

$$|\psi_{\mathbf{s}}\rangle = e^{-i\theta_{\mathbf{s}}}\,|\mathbf{s}_{\mathcal{C}'}\rangle\,, \tag{E78}$$

where $\{|\mathbf{s}_{\mathcal{C}'}\rangle\}$ is a logical basis for an $[[n,k]]$ CSS code $\mathcal{C}'$ generated by a completely CSS-preserving unitary.

Proof.—By Eq. (B22), the stabilizer group for the logical basis state $|\mathbf{s}_{\mathcal{C}}\rangle$ can be related to the stabilizer group $\mathcal{S}(\mathcal{C})$ of $\mathcal{C}$ as

$$\mathcal{S}(|\mathbf{s}_{\mathcal{C}}\rangle) = \langle(-1)^{s_1}S_1,\ldots,(-1)^{s_k}S_k\rangle\times\mathcal{S}(\mathcal{C}), \tag{E79}$$

where we define $S_i := UZ_iU^{\dagger}$ for $i=1,\ldots,k$.

We observe that $\mathcal{S}(\mathcal{C})$ is a direct sum of $\mathbb{F}_2$ subspaces $\mathcal{S}_X(\mathcal{C})$ and $\mathcal{S}_Z(\mathcal{C})$ corresponding to X- and Z-type stabilizers for $\mathcal{C}$. Because $|\psi\rangle$ is a CSS state, its stabilizer group $\mathcal{S}(|\psi\rangle)$ is similarly a direct sum of $\mathbb{F}_2$ subspaces $\mathcal{S}_Z(|\psi\rangle)$ and $\mathcal{S}_X(|\psi\rangle)$ corresponding to X- and Z-type stabilizers for $|\psi\rangle$.

Since $\{|\mathbf{s}_{\mathcal{C}}\rangle\}$ span $\mathcal{C}$, Eq. (E77) implies that, if S stabilizes $|\psi\rangle$, then $\mathbb{1}^{\otimes n}\otimes S$ stabilizes $\mathcal{C}$. Therefore, $\mathbb{1}^{\otimes n}\otimes\mathcal{S}_Z(|\psi\rangle)$ and $\mathbb{1}^{\otimes n}\otimes\mathcal{S}_X(|\psi\rangle)$ are $\mathbb{F}_2$ subspaces of $\mathcal{S}_Z(\mathcal{C})$ and $\mathcal{S}_X(\mathcal{C})$, respectively. By applying the basis extension theorem separately to the Z and X cases, we can represent $\mathcal{S}(\mathcal{C})$ as

$$\mathcal{S}(\mathcal{C}) = \langle S_{k+1},\ldots,S_n,\mathbb{1}^{\otimes n}\otimes T_1,\ldots,\mathbb{1}^{\otimes n}\otimes T_m\rangle, \tag{E80}$$

where $S_{k+1},\ldots,S_n$ and $T_1,\ldots,T_m$ are all CSS observables such that $\mathcal{S}(|\psi\rangle) = \langle T_1,\ldots,T_m\rangle$.

The stabilizer group of $|\mathbf{s}_{\mathcal{C}}\rangle$ can then be represented as

$$\begin{aligned}\mathcal{S}(|\mathbf{s}_{\mathcal{C}}\rangle) = \langle(-1)^{s_1}S_1,\ldots,(-1)^{s_k}S_k,S_{k+1},\ldots,S_n,\\ \mathbb{1}^{\otimes n}\otimes T_1,\ldots,\mathbb{1}^{\otimes n}\otimes T_m\rangle,\end{aligned} \tag{E81}$$

from which we note that $S_1,\ldots,S_n$ are all members of $\mathcal{S}(|\mathbf{0}_{\mathcal{C}}\rangle)$, the stabilizer group for the zero logical basis state in $\mathcal{C}$. Since $|\mathbf{0}_{\mathcal{C}}\rangle$ and $|\psi\rangle$ are CSS [see Eq. (B22)], $|\psi_{\mathbf{0}}\rangle$ must be CSS as well (as discarding the last m qubits is CSS preserving). Therefore, $\mathcal{S}(|\mathbf{0}_{\mathcal{C}}\rangle)$ is a direct sum of $\mathbb{F}_2$ subspaces $\mathbb{1}^{\otimes n}\otimes\mathcal{S}(|\psi\rangle)$ and $\mathcal{S}(|\psi_{\mathbf{0}}\rangle)\otimes\mathbb{1}^{\otimes m}$, where $\mathcal{S}(|\psi_{\mathbf{0}}\rangle)$ is the stabilizer group of $|\psi_{\mathbf{0}}\rangle$. Consequently, we can express S_i as $S_i = S_i'\otimes T_i'$, where S_i' and T_i' are CSS observables of the same type as S_i that, respectively, stabilize $|\psi_{\mathbf{0}}\rangle$ and $|\psi\rangle$. Therefore, by multiplying each S_i in Eq. (E81) by appropriate generators of the same type in $\{\mathbb{1}^{\otimes n}\otimes T_1,\ldots\mathbb{1}^{\otimes n}\otimes T_m\}$, we can represent the stabilizer group of $|\mathbf{s}_{\mathcal{C}}\rangle$ as

$$\mathcal{S}(|\mathbf{s}_{\mathcal{C}}\rangle) = \mathcal{S}(|\psi_{\mathbf{s}}\rangle)\otimes\mathcal{S}(|\psi\rangle), \tag{E82}$$

where the stabilizer group $\mathcal{S}(|\psi_{\mathbf{s}}\rangle)$ is generated by

$$\mathcal{S}(|\psi_{\mathbf{s}}\rangle) = \langle(-1)^{s_1}S_1',\ldots,(-1)^{s_k}S_k',S_{k+1}',\ldots,S_n'\rangle, \tag{E83}$$

in which $S_1',\ldots,S_k'$ are Z type because $S_1,\ldots,S_k$ are Z type. By applying Lemma 9 to $S_1',\ldots,S_n'$, we can find a logical basis $\{|\mathbf{s}_{\mathcal{C}'}\rangle\}$ for an $[[n,k]]$ CSS code $\mathcal{C}'$ stabilized by $\langle S_{k+1}',\ldots,S_n'\rangle$, generated by a completely CSS-preserving encoding unitary, such that $|\mathbf{s}_{\mathcal{C}'}\rangle$ shares the stabilizer group of $|\psi_{\mathbf{s}}\rangle$. Therefore, $|\mathbf{s}_{\mathcal{C}'}\rangle$ and $|\psi_{\mathbf{s}}\rangle$ differ only up to a phase, which implies the lemma. ■

Lemma 21.—Let $\mathcal{K}$ be the code projection for an $[[n+m,k]]$ CSS code $\mathcal{C}$ where $n\geq 1, n>k$ and $m>0$. Then given any m-qubit CSS state $|\psi\rangle$, we have that $\mathcal{C}([\cdot]\otimes|\psi\rangle\langle\psi|)$ is equivalent to a CSS-code projection on n qubits *alone*, followed by preparing a CSS state and completely CSS-preserving postprocessing, i.e.,

$$\mathcal{K}(\rho\otimes|\psi\rangle\langle\psi|) = p\,\tilde{\mathcal{U}}\circ\left(\tilde{\mathcal{K}}(\rho)\otimes|\varphi\rangle\langle\varphi|\right), \tag{E84}$$

where p is a probability, $\tilde{\mathcal{U}}$ is a completely CSS-preserving unitary channel on k qubits, $\tilde{\mathcal{K}}$ is a code projection for an $[[n,k']]$ CSS code where $0\leq k'\leq k$, and $|\varphi\rangle$ is a CSS state on $k-k'$ qubits.

Proof.—Let $\{S_{n+1},\ldots,S_{n+m}\}$ be a set of CSS observables that generate the stabilizer group defining $|\psi\rangle$. Then $\mathcal{K}(\rho\otimes|\psi\rangle\langle\psi|)$ is equivalent to

$$\mathcal{K}(\rho\otimes|\psi\rangle\langle\psi|) = \mathcal{K}(\mathbf{P}[\rho\otimes|\psi\rangle\langle\psi|]\mathbf{P}), \tag{E85}$$

where $\mathbf{P}$ projects the last m qubits onto $|\psi\rangle$, i.e.,

$$\mathbf{P} := \mathbb{1}^{\otimes n}\otimes|\psi\rangle\langle\psi| = \prod_{i=n+1}^{n+m} P(\mathbb{1}^{\otimes n}\otimes S_i). \tag{E86}$$

By applying Lemma 19 to each projection carried out by $\mathbf{P}$, we obtain

$$\mathcal{K}(\rho\otimes|\psi\rangle\langle\psi|) = p\,\mathcal{U}\circ\mathcal{K}'(\rho\otimes|\psi\rangle\langle\psi|)\otimes|\varphi\rangle\langle\varphi|, \tag{E87}$$

where p is a probability, $\mathcal{U}$ is a completely CSS-preserving unitary channel on k qubits, $\mathcal{K}'$ is a code projection for an $[[n+m,k']]$ CSS code $\mathcal{C}'$ where $0\leq k'\leq k$, and $|\varphi\rangle$ is a CSS state on $k-k'$ qubits. Lemma 19 further implies there exists a logical basis $\{|\mathbf{s}_{\mathcal{C}'}\rangle\,|\,\mathbf{s}\in\{0,1\}^{k'}\}$ for the new code $\mathcal{C}'$, generated by a completely CSS-preserving encoding unitary, which can be related to some state $|\Psi_{\mathbf{s}}\rangle$ on $n+m$ qubits as

$$\mathbf{P}\,|\Psi_{\mathbf{s}}\rangle = \left(\mathbb{1}^{\otimes n}\otimes|\psi\rangle\langle\psi|\right)|\Psi_{\mathbf{s}}\rangle \propto |\mathbf{s}_{\mathcal{C}'}\rangle\,. \tag{E88}$$

Equation Eq. (E88) immediately implies

$$\left|\mathbf{s}'_{\mathcal{C}'}\right\rangle = |\psi_{\mathbf{s}}\rangle\otimes|\psi\rangle\,, \tag{E89}$$

where $|\psi_{\mathbf{s}}\rangle$ is a state on the first n qubits *alone*. By Lemma 20, each $|\psi_{\mathbf{s}}\rangle$ is, *up to a phase that may vary with* $\mathbf{s}$, a logical basis state $|\mathbf{s}_{\mathcal{C}''}\rangle$ for an $[[n,k']]$ CSS code $\mathcal{C}''$ with only n physical qubits. Thus one can write

$$|\mathbf{s}_{\mathcal{C}'}\rangle = e^{-i\theta_{\mathbf{s}}}\,|\mathbf{s}_{\mathcal{C}''}\rangle\otimes|\psi\rangle\,. \tag{E90}$$

By Eq. (E39), we can express the code projection $\mathcal{K}'$ for $\mathcal{C}'$ in terms of the logical basis $\{|\mathbf{s}_{\mathcal{C}'}\rangle\}$ as

$$\mathcal{K}'(\cdot) = K(\cdot)K^{\dagger},\; K^{\dagger} := \sum_{\mathbf{s}\in\{0,1\}^{k'}} |\mathbf{s}_{\mathcal{C}'}\rangle\langle\mathbf{s}|\,. \tag{E91}$$

We can then use Eq. (E90) to show that

$$\mathcal{K}'(\rho\otimes|\psi\rangle\langle\psi|) = U'\left[\tilde{\mathcal{K}}(\rho)\right]U'^{\dagger}, \tag{E92}$$

where we define the following unitary on k' qubits to adjust for the phase differences between $|\mathbf{s}_{\mathcal{C}'}\rangle$ and $|\psi_{\mathbf{s}}\rangle$,

$$U' := \sum_{\mathbf{s}\in\{0,1\}^{k'}} e^{i\theta_{\mathbf{s}}}\,|\mathbf{s}\rangle\langle\mathbf{s}|\,, \tag{E93}$$

as well as

$$\tilde{\mathcal{K}}(\cdot) := \tilde{K}(\cdot)\tilde{K}^{\dagger},\; \tilde{K}^{\dagger} := \sum_{\mathbf{s}\in\{0,1\}^{k'}} |\mathbf{s}_{\mathcal{C}''}\rangle\langle\mathbf{s}|\,, \tag{E94}$$

By comparison with Eq. (E39), we see that $\tilde{\mathcal{K}}$ is the code projection for the $[[n,k']]$ CSS code $\mathcal{C}''$. Following a similar argument to that at the end of case (iii) in the proof of Lemma 19, we can demonstrate that U' is completely CSS preserving. Substituting back immediately yields the lemma result. ■

APPENDIX F: COMPLEX RELATIVE MAJORIZATION

In this Appendix, we sketch how our entropic constraints can be extended to arbitrary input and output states—i.e., those whose density matrices in the computational basis may not be real.

We refer to any complex-valued distribution $\mathbf{w}$ on N elements as a *complex quasidistribution* if its components sum to 1, i.e.,

$$\sum_{i=1}^{N} \mathbf{w}^{(i)} = 1. \tag{F1}$$

Given complex quasidistributions $\mathbf{w}, \mathbf{w}'$ and reference probability distributions $\mathbf{r}, \mathbf{r}'$, we can extend relative majorization very naturally to a partial order $\succ_C$ defined by

$$(\mathbf{w}, \mathbf{r}) \succ_C (\mathbf{w}', \mathbf{r}') \iff A\mathbf{w} = \mathbf{w}',\ A\mathbf{r} = \mathbf{r}'. \tag{F2}$$

Any complex quasidistribution $\mathbf{w}$ on N elements can be decomposed into real and imaginary parts as

$$\mathbf{w} = \mathbf{w}_R + i\mathbf{w}_I, \tag{F3}$$

where we introduce the notation $(\cdot)_R := \mathrm{Re}(\cdot)$ and $(\cdot)_I := \mathrm{Im}(\cdot)$. By Eq. (F1), we have

$$\sum_{i=1}^{N} w_R^{(i)} = 1,\ \sum_{i=1}^{N} w_I^{(i)} = 0. \tag{F4}$$

Therefore, given any complex quasidistribution $\mathbf{w}$ on N elements, we can construct a valid *quasiprobability distribution* on $2N$ elements using the map

$$\mu(\mathbf{w}) := \mathbf{w}_R \oplus \mathbf{w}_I. \tag{F5}$$

Furthermore, given any probability distribution $\mathbf{r}$ on N elements, we can also construct a probability distribution on $2N$ elements using the map

$$\nu(\mathbf{r}) := \frac{1}{2}(\mathbf{r} \oplus \mathbf{r}). \tag{F6}$$

We can now prove the following theorem.

Theorem 7.—Given complex quasidistributions $\mathbf{w}, \mathbf{w}'$ and probability distributions $\mathbf{r}, \mathbf{r}'$, we have that

$$(\mathbf{w}, \mathbf{r}) \succ_C (\mathbf{w}', \mathbf{r}') \implies (\mu(\mathbf{w}), \nu(\mathbf{r})) \succ (\mu(\mathbf{w}'), \nu(\mathbf{r}')). \tag{F7}$$

Proof.—Since a stochastic matrix processes the real and imaginary parts of a vector independently, we conclude that $(\mathbf{w}, \mathbf{r}) \succ_C (\mathbf{w}', \mathbf{r}')$ implies

$$\begin{bmatrix} \mathbf{w}'_R \\ \mathbf{w}'_I \end{bmatrix} = \begin{bmatrix} A & 0 \\ 0 & A \end{bmatrix} \begin{bmatrix} \mathbf{w}_R \\ \mathbf{w}_I \end{bmatrix}, \begin{bmatrix} \frac{1}{2}\mathbf{r}' \\ \frac{1}{2}\mathbf{r}' \end{bmatrix} = \begin{bmatrix} A & 0 \\ 0 & A \end{bmatrix} \begin{bmatrix} \frac{1}{2}\mathbf{r} \\ \frac{1}{2}\mathbf{r} \end{bmatrix} \tag{F8}$$

for some stochastic matrix A. Defining $\tilde{A} := A \oplus A$, we can rewrite this more compactly as

$$\mu(\mathbf{w}') = \tilde{A}\mu(\mathbf{w}),\ \nu(\mathbf{r}') = \tilde{A}\nu(\mathbf{r}), \tag{F9}$$

where $\tilde{A}$ is stochastic whenever A is stochastic, completing the proof. ■

Corollary 5.—Let $\mathbf{w}_\rho := \mathrm{Re}[W_\rho] \oplus \mathrm{Im}[W_\rho]$ and $\mathbf{r}_\tau := 1/2(W_\tau \oplus W_\tau)$. If there exists a completely CSS-preserving channel $\mathcal{E}$ such that $\mathcal{E}(\rho) = \rho'$ and $\mathcal{E}(\tau) = \tau'$, where τ and τ' are both in the interior of $\mathcal{D}_{\mathrm{CSS}}$, then for all $\alpha \in \mathcal{A}$,

$$D_\alpha(\mathbf{w}_\rho || \mathbf{r}_\tau) \geq D_\alpha(\mathbf{w}_{\rho'} || \mathbf{r}_{\tau'}). \tag{F10}$$

[1] E. T. Campbell, B. M. Terhal, and C. Vuillot, Roads towards fault-tolerant universal quantum computation, Nature **549**, 172 (2017).

[2] R. Raussendorf and H. J. Briegel, A One-Way Quantum Computer, Phys. Rev. Lett. **86**, 5188 (2001).

[3] R. Raussendorf, D. E. Browne, and H. J. Briegel, Measurement-based quantum computation on cluster states, Phys. Rev. A **68**, 022312 (2003).

[4] N. H. Nickerson, J. F. Fitzsimons, and S. C. Benjamin, Freely Scalable Quantum Technologies Using Cells of 5-to-50 Qubits with Very Lossy and Noisy Photonic Links, Phys. Rev. X **4**, 041041 (2014).

[5] E. Nikahd, M. Sedighi, and M. Saheb Zamani, Nonuniform code concatenation for universal fault-tolerant quantum computing, Phys. Rev. A **96**, 032337 (2017).

[6] R. Chao and B. W. Reichardt, Fault-tolerant quantum computation with few qubits, Npj Quantum Inf. **4**, 1 (2018).

[7] C. Lin, G. Yang, Q. Luo, and X. Li, Pieceable fault tolerant conversion between 5-qubit code and 7-CSS code, Quantum Inf. Process. **19**, 243 (2020).

[8] J. E. Bourassa, R. N. Alexander, M. Vasmer, A. Patil, I. Tzitrin, T. Matsuura, D. Su, B. Q. Baragiola, S. Guha, G. Dauphinais, K. K. Sabapathy, N. C. Menicucci, and I. Dhand, Blueprint for a scalable photonic fault-tolerant quantum computer, Quantum **5**, 392 (2021).

[9] C. Chamberland, K. Noh, P. Arrangoiz-Arriola, E. T. Campbell, C. T. Hann, J. Iverson, H. Putterman, T. C. Bohdanowicz, S. T. Flammia, A. Keller, G. Refael, J. Preskill, L. Jiang, A. H. Safavi-Naeini, O. Painter, and F. G. Brandão, Building a Fault-Tolerant Quantum Computer using Concatenated Cat Codes, PRX Quantum **3**, 010329 (2022).

[10] S. B. Bravyi and A. Y. Kitaev, Quantum codes on a lattice with boundary, arXiv:quant-ph/9811052 (1998).

[11] E. Dennis, A. Kitaev, A. Landahl, and J. Preskill, Topological quantum memory, J. Math. Phys. **43**, 4452 (2002).

[12] M. H. Freedman, Quantum computation and the localization of modular functors, Found. Comput. Math. **1**, 183 (2001).

[13] D. Gottesman, Ph.D. thesis, Caltech, 1997.

[14] D. Gottesman, The Heisenberg representation of quantum computers, arXiv:quant-ph/9807006 (1998).

[15] B. Eastin and E. Knill, Restrictions on Transversal Encoded Quantum Gate Sets, Phys. Rev. Lett. **102**, 110502 (2009).

[16] V. Veitch, C. Ferrie, D. Gross, and J. Emerson, Negative quasi-probability as a resource for quantum computation, New J. Phys. **14**, 113011 (2012).

[17] 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).

[18] S. Bravyi and A. Kitaev, Universal quantum computation with ideal Clifford gates and noisy ancillas, Phys. Rev. A **71**, 022316 (2005).

[19] B. W. Reichardt, Quantum universality from magic states distillation applied to CSS codes, Quantum Inf. Process. **4**, 251 (2005).

[20] S. Bravyi and J. Haah, Magic-state distillation with low overhead, Phys. Rev. A **86**, 052329 (2012).

[21] J. Haah, M. B. Hastings, D. Poulin, and D. Wecker, Magic state distillation with low space overhead and optimal asymptotic input count, Quantum **1**, 31 (2017).

[22] M. B. Hastings and J. Haah, Distillation with Sublogarithmic Overhead, Phys. Rev. Lett. **120**, 050504 (2018).

[23] A. R. Calderbank and P. W. Shor, Good quantum error-correcting codes exist, Phys. Rev. A **54**, 1098 (1996).

[24] A. Steane, Multiple-particle interference and quantum error correction, Proc. R. Soc. Lond. A **452**, 2551 (1996).

[25] N. Rengaswamy, R. Calderbank, M. Newman, and H. D. Pfister, On optimality of CSS codes for transversal t, IEEE J. Sel. Areas Inf. Theory **1**, 499 (2020).

[26] D. Litinski, Magic state distillation: Not as costly as you think, Quantum **3**, 205 (2019).

[27] N. Koukoulekidis and D. Jennings, Constraints on magic state protocols from the statistical mechanics of Wigner negativity, Npj Quantum Inf. **8**, 42 (2022).

[28] D. Gross, Hudson's theorem for finite-dimensional quantum systems, J. Math. Phys. **47**, 122107 (2006).

[29] D. M. Appleby, Symmetric informationally complete–positive operator valued measures and the extended Clifford group, J. Math. Phys. **46**, 052107 (2005).

[30] A. F. Veinott, Least *d*-majorized network flows with inventory and statistical applications, Manag. Sci. **17**, 547 (1971).

[31] M. Horodecki and J. Oppenheim, Fundamental limitations for quantum and nanoscale thermodynamics, Nat. Commun. **4**, 1 (2013).

[32] D. Blackwell, Equivalent comparisons of experiments, Ann. Math. Stat **24**, 265 (1953).

[33] E. Ruch and A. Mead, The principle of increasing mixing character and some of its consequences, Theor. Chim. Acta **41**, 95 (1976).

[34] M. Lostaglio, An introductory review of the resource theory approach to thermodynamics, Rep. Prog. Phys **82**, 114001 (2019).

[35] D. Gross, Ph.D. thesis, Imperial College London, 2005.

[36] A. Mari and J. Eisert, Positive Wigner Functions Render Classical Simulation of Quantum Computation Efficient, Phys. Rev. Lett. **109**, 230503 (2012).

[37] E. F. Galvão, Discrete Wigner functions and quantum computational speedup, Phys. Rev. A **71**, 042302 (2005).

[38] C. Cormick, E. F. Galvão, D. Gottesman, J. P. Paz, and A. O. Pittenger, Classicality in discrete Wigner functions, Phys. Rev. A **73**, 012301 (2006).

[39] V. Veitch, C. Ferrie, D. Gross, and J. Emerson, Negative quasi-probability as a resource for quantum computation, New J. Phys. **14**, 113011 (2012).

[40] H. Anwar, E. T. Campbell, and D. E. Browne, Qutrit magic state distillation, New J. Phys. **14**, 063006 (2012).

[41] R. Raussendorf, J. Bermejo-Vega, E. Tyhurst, C. Okay, and M. Zurel, Phase-space-simulation method for quantum computation with magic states on qubits, Phys. Rev. A **101**, 012350 (2020).

[42] M. Zurel, C. Okay, and R. Raussendorf, Hidden Variable Model for Universal Quantum Computation with Magic States on Qubits, Phys. Rev. Lett. **125**, 260404 (2020).

[43] L. Catani and D. E. Browne, Spekkens' toy model in all dimensions and its relationship with stabiliser quantum mechanics, New J. Phys. **19**, 073035 (2017).

[44] N. Delfosse, P. Allard Guerin, J. Bian, and R. Raussendorf, Wigner Function Negativity and Contextuality in Quantum Computation on Rebits, Phys. Rev. X **5**, 021003 (2015).

[45] L. Catani and D. E. Browne, State-injection schemes of quantum computation in Spekkens' toy theory, Phys. Rev. A **98**, 052108 (2018).

[46] B. Regula, Probabilistic Transformations of Quantum Resources, Phys. Rev. Lett. **128**, 110505 (2022).

[47] E. T. Campbell and D. E. Browne, in *Theory of Quantum Computation, Communication, and Cryptography*, edited by A. Childs and M. Mosca (Springer Berlin Heidelberg, Berlin, Heidelberg, 2009), p. 20.

[48] K. Fang and Z.-W. Liu, No-Go Theorems for Quantum Resource Purification, Phys. Rev. Lett. **125**, 060405 (2020).

[49] I. Bengtsson and K. Życzkowski, *Geometry of Quantum States: An Introduction to Quantum Entanglement* (Cambridge University Press, Cambridge, 2006).

[50] S. Bravyi, G. Smith, and J. A. Smolin, Trading Classical and Quantum Computational Resources, Phys. Rev. X **6**, 021043 (2016).

[51] J. Watrous, *The Theory of Quantum Information* (Cambridge University Press, Cambridge, 2018).

[52] X. Wang, M. M. Wilde, and Y. Su, Quantifying the magic of quantum channels, New J. Phys. **21**, 103002 (2019).

[53] D. Schmid, H. Du, J. H. Selby, and M. F. Pusey, Uniqueness of Noncontextual Models for Stabilizer Subtheories, Phys. Rev. Lett. **129**, 120403 (2022).

[54] P. Allard Guérin, Master's thesis, The University of British Columbia, 2015.

[55] R. Raussendorf and J. Harrington, Fault-Tolerant Quantum Computation with High Threshold in Two Dimensions, Phys. Rev. Lett. **98**, 190504 (2007).

[56] V. Veitch, S. A. H. Mousavian, D. Gottesman, and J. Emerson, The resource theory of stabilizer quantum computation, New J. Phys. **16**, 013009 (2014).

[57] M. Howard and E. Campbell, Application of a Resource Theory for Magic States to Fault-Tolerant Quantum Computing, Phys. Rev. Lett. **118**, 090501 (2017).

[58] X. Wang, M. M. Wilde, and Y. Su, Efficiently Computable Bounds for Magic State Distillation, Phys. Rev. Lett. **124**, 090505 (2020).

[59] K. Fang and Z.-W. Liu, No-Go Theorems for Quantum Resource Purification: New Approach and Channel Theory, PRX Quantum **3**, 010337 (2022).

[60] 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).

[61] Note that since the set $\mathcal{A}$ is dense in $\alpha \in (1, \infty)$, one could extend the definition of $D_\alpha(\mathbf{w}||\mathbf{r})$ to all $\alpha \geq 1$ by continuity.

[62] We note that an analogous set of magic monotones can be defined for systems of odd-prime dimension via

$$\Gamma_\alpha(\rho) = \inf_{\tau \in STAB} D_\alpha(W_\rho||W_\tau). \tag{F11}$$

.

[63] J. R. Seddon, B. Regula, H. Pashayan, Y. Ouyang, and E. T. Campbell, Quantifying Quantum Speedups: Improved Classical Simulation from Tighter Magic Monotones, PRX Quantum **2**, 010345 (2021).

[64] E. Knill, R. Laflamme, and W. Żurek, Threshold accuracy for quantum computation, arXiv:quant-ph/9610011 (1996).

[65] A. M. Steane, Quantum Reed-Muller codes, IEEE Trans. Inf. Theory **45**, 1701 (1999).

[66] Under the restriction to the rebit subtheory of quantum computing subject to CSS circuits, the link between negativity and magic has been restored [44]. It, therefore, follows that mana is a monotone under the class of distillation protocols considered here. In fact, the monotonicity of mana is equivalent to $\Delta D_\alpha \geq 0$ in the limit $\alpha \to 1$.

[67] S. Beigi, Sandwiched rényi divergence satisfies data processing inequality, J. Math. Phys. **54**, 122202 (2013).

[68] M. Mosonyi, in *9th Conference on the Theory of Quantum Computation, Communication and Cryptography, May 21-23, 2014, Singapore*, LIPIcs, Vol. 27, edited by S. T. Flammia and A. W. Harrow (Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2014), p. 88.

[69] M. M. Wilde, A. Winter, and D. Yang, Strong converse for the classical capacity of entanglement-breaking and Hadamard channels via a sandwiched Rényi relative entropy, Commun. Math. Phys. **331**, 593 (2014).

[70] A. Kitaev, Fault-tolerant quantum computation by anyons, Ann. Phys. **303**, 2 (2003).

[71] F. Albarelli, M. G. Genoni, M. G. A. Paris, and A. Ferraro, Resource theory of quantum non-gaussianity and Wigner negativity, Phys. Rev. A **98**, 052350 (2018).

[72] N. Moiseyev, *Non-Hermitian Quantum Mechanics* (Cambridge University Press, Cambridge, 2011).

[73] A. Giovagnoli and H. P. Wynn, Cyclic majorization and smoothing operators, Linear Algebra Appl. **239**, 215 (1996).

[74] A. Giovagnoli and H. P. Wynn, G-majorization with applications to matrix orderings, Linear Algebra Appl. **67**, 111 (1985).

[75] A. G. M. Steerneman, G-majorization, group-induced cone orderings, and reflection groups, Linear Algebra Appl. **127**, 107 (1990).

[76] R. Bhatia, *Matrix Analysis* (Springer Science & Business Media, New York, 2013), Vol. 169.

[77] M. P. Woods and M. Horodecki, The resource theoretic paradigm of quantum thermodynamics with control, arXiv:quant-ph/1912.05562 (2019).