Hardness of recognizing phases of matter

Thomas Schuster[*,1,2], Dominik Kufel[*,3], Norman Y. Yao[3], and Hsin-Yuan Huang[1,2]

[1]California Institute of Technology
[2]Google Quantum AI
[3]Harvard University

October 10, 2025

Abstract

We prove that recognizing the phase of matter of an unknown quantum state is quantum computationally hard. More specifically, we show that the quantum computational time of any phase recognition algorithm must grow exponentially in the range of correlations ξ of the unknown state. This exponential growth renders the problem practically infeasible for even moderate correlation ranges, and leads to super-polynomial quantum computational time in the system size n whenever $\xi = \omega(\log n)$. Our results apply to a substantial portion of all known phases of matter, including symmetry-breaking phases and symmetry-protected topological phases for any discrete on-site symmetry group in any spatial dimension. To establish this hardness, we extend the study of pseudorandom unitaries (PRUs) to quantum systems with symmetries. We prove that symmetric PRUs exist under standard cryptographic conjectures, and can be constructed in extremely low circuit depths. We also establish hardness for systems with translation invariance and purely classical phases of matter. A key technical limitation is that the locality of the parent Hamiltonians of the states we consider is linear in ξ; the complexity of phase recognition for Hamiltonians with constant locality remains an important open question.

1 Introduction

Phases of matter form a cornerstone of modern physics, and provide a fundamental organizing principle for understanding complex quantum systems. Classical phases of matter are famously characterized by the Landau paradigm, where phase transitions are driven by the spontaneous breaking of symmetries in the system [1]. Beyond the Landau paradigm, quantum mechanics has unveiled entirely new phases of matter, including topological order and symmetry-protected topological phases [2–4]. The ability to identify and characterize these diverse phases of matter is of fundamental interest across physics and information science and crucial for advancing quantum technologies [5–16].

A fundamental question underlying this scientific pursuit is the following: Given experimental access to a quantum state, how can one determine its phase of matter [17–26]? Despite tremendous interest, answering this question has proven challenging. All rigorous algorithms for phase recognition have runtimes that scale exponentially in the range of correlations ξ of the state [27–29]. This scaling mirrors that of brute-force quantum state tomography on local patches of radius $\mathcal{O}(\xi)$ and becomes infeasible for physical systems with even a moderate correlation range. Intriguingly, recent work

has proven that for certain phases, this computational hardness is fundamental: recognizing whether a quantum state has *topological order* requires super-polynomial quantum computational resources whenever $\xi = \omega(\log n)$ [30]. This follows from the discovery that pseudorandom unitaries (PRUs) [31–34] can be realized in extremely shallow geometrically-local circuits [30]. However, this result leaves open an essential question: *How hard is it to recognize phases of matter beyond topological order?*

In this work, we prove that a substantial portion of all known phases of matter are quantum computationally hard to recognize. This includes symmetry-breaking phases and symmetry-protected topological phases for any discrete on-site symmetry group in any spatial dimension. Our results apply equally to ground states and mixed states, and encompass phases both with and without translation invariance (over sufficiently large unit cells). We also extend our results to purely *classical* phases of matter using distinct techniques. In all of these settings, we prove that the complexity of recognizing the phase of matter of an unknown quantum or classical state grows exponentially in the range of correlations ξ of the state[1]. Asymptotically, the quantum computational complexity becomes super-polynomial in the system size n whenever $\xi = \omega(\log n)$.

Our results apply to any definition of phases of matter that is invariant under shallow symmetric geometrically-local unitary circuits. An important technical limitation is that the locality of the parent Hamiltonian of the states we consider must scale linearly in the range of correlations ξ. This contrasts with typical physical settings, in which Hamiltonians are 2-local. Addressing the complexity of phase recognition for such Hamiltonians remains an important open question for future work.

To establish these results, we extend the notion of PRUs, and their formation in extremely low circuit depths, to incorporate a broad range of physical symmetry operations. A PRU is an efficient quantum circuit that is indistinguishable from a completely random unitary by any polynomial-time quantum computation [31, 34, 35]. Together with their close cousin, unitary designs [30, 36–45], PRUs are central to numerous areas of quantum science, including quantum cryptography [31, 46, 47], device benchmarking [48–50], quantum supremacy experiments [51–53], quantum learning theory [54–58], quantum gravity [59–63], and many-body quantum dynamics [64–69]. A critical question in recent research is how to generalize PRUs to physically constrained classes of quantum dynamics, such as those with symmetries [30, 70–79]. Several works have established no-go results, showing that low-depth constructions of PRUs are impossible in the presence of continuous [77] or time-reversal symmetries [30, 78, 79]. To achieve our results, we develop an extensive collection of tools to prove that nearly all features of PRUs *do* extend to quantum systems with any discrete on-site symmetries. Beyond their implications for recognizing phases of matter, our results thus provide broader support for using PRUs as models of realistic quantum systems in the physical world.

At a conceptual level, our results should be viewed as a *worst-case* statement: There exist classical and quantum states whose phase of matter is precisely defined, yet is impossible to recognize in any efficient quantum experiment. This statement feels counterintuitive at first impression. Daily experience teaches us that symmetry-breaking phases of matter, such as liquids, solids, and magnets, are easy to recognize in practical scenarios. Recent work has even argued that this easiness extends to symmetry-protected topological phases, using ideas from the renormalization group [24]. Nonetheless, we rigorously prove that there exist classical and quantum states where all approaches to phase recognition must fail. Looking forward, this raises a critical open question: If not symmetries, what properties of quantum states in the natural world guarantee that their phases of matter are so easily recognized? Answering this question will require new insights into quantum phases of matter beyond our existing knowledge based on spectral gaps and Lieb-Robinson bounds.

[1] Here, the range of correlations ξ refers to the maximum distance at which any connected correlation functions remain non-zero or above a small value; the state may exhibit fully nonlocal correlations within distance ξ. This differs slightly from other notions of correlation length, which consider the asymptotic exponential decay rate of correlation functions. Our results show that the complexity of phase recognition grows exponentially in the range of correlations ξ, independent of the asymptotic exponential decay rate.

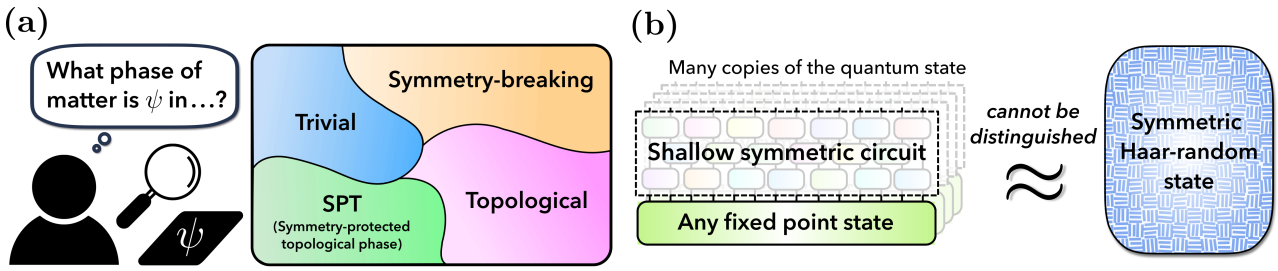

Figure 1: Illustration of our main results. **(a)** We consider the question: Given experimental access to many copies of a quantum state $|\psi\rangle$, can one recognize what phase of matter $|\psi\rangle$ is in? Examples of phases of matter include the trivial phase, symmetry-breaking phases, symmetry-protected topological (SPT) phases, and topological order. Our main result is that the complexity of recognizing phases of matter grows exponentially in the correlation range ξ of the state, becoming super-polynomial in the system size n as soon as $\xi = \text{poly}(\log n)$. **(b)** We achieve this result by extending the study of pseudorandom unitaries (PRUs) to quantum systems with symmetries. We show that any fixed point state of any phase of matter can become indistinguishable from a symmetric Haar-random state after a low-depth symmetric circuit is applied.

The remainder of our manuscript is structured as follows. Section 2.1 provides background on quantum phases of matter, and Section 2.2 summarizes our main results on recognizing quantum phases of matter. These results are founded upon our low-depth construction of symmetric PRUs, which we describe in Sections 2.3 and 2.4. We return to quantum phases of matter in Section 2.5 to provide clarifying discussions, intuition, and detailed analyses that demonstrate how hardness can emerge even at small correlation lengths for common phase recognition approaches. Section 2.6 extends our results to mixed state quantum phases of matter. We then present our results on classical phases of matter in Section 3. Section 4 concludes with additional discussions and open questions.

2 Quantum phases of matter

In this section, we describe our main results on recognizing phases of matter in quantum systems.

2.1 Background

Quantum phases of matter capture properties of the ground states of gapped many-body quantum systems that remain invariant under smooth deformations of the system [2, 4]. The simplest example is a symmetry-breaking phase, in which the system is invariant under a symmetry G, yet features multiple nearly-degenerate ground states that transform into one another under the action of G [1]. This contrasts with the trivial phase of a symmetric system, which features a unique ground state that respects the symmetry operation. Quantum states can also exhibit novel phases beyond the symmetry-breaking paradigm. These include topological orders [2, 80], which can occur in the absence of any symmetry and are characterized by their patterns of long-range entanglement, and symmetry-protected topological (SPT) phases [3, 4], which are short-range entangled, yet in a manner that is non-trivial when the symmetry is considered. We refer to Section 2.5 for several concrete examples.

Our modern understanding of how to classify phases of matter centers upon ideas from quantum information [4]. Two quantum states are in the same phase of matter if and only if they are connected by (i) a low-depth unitary circuit, or (ii) a short-time adiabatic evolution under gapped Hamiltonians [81–87]. The two definitions are equivalent up to moderate overheads: any short-time adiabatic evolution can be Trotterized to a low-depth unitary circuit up to a $\text{poly}(\log n)$ overhead in

circuit depth, while any depth-ℓ unitary circuit yields a gapped adiabatic path interpolating between an original parent Hamiltonian, H_1, and a final parent Hamiltonian, $H_2 = UH_1U^\dagger$, up to a $\mathrm{poly}(\ell)$ increase in the Hamiltonian's locality. Intuitively, low-depth circuits and short-time evolutions cannot change the structure of correlations and entanglement in a quantum state, owing to their limited light-cone. Hence, the phase of matter remains invariant. If the system is symmetric, the circuit and Hamiltonian must respect the symmetry operation as well.

Our results apply to any definition of phases of matter such that the following holds. For each phase of matter, we can associate a representative "fixed point" state $|\psi_0\rangle$. Then, we assume that any quantum state $|\psi\rangle$ that is related to $|\psi_0\rangle$ by a symmetric geometrically-local depth-ℓ unitary circuit U is in the same phase as $|\psi_0\rangle$. The parameter ℓ controls the maximum allowed circuit depth of the unitary U. We also let ξ denote the radius of the light-cone of U. Physically, ξ upper bounds the maximum range of correlations in the quantum state $|\psi\rangle$. We have $\xi \leq \ell$ by definition.

Oftentimes in the literature, the correlation range ξ and circuit depth ℓ are assumed to be constant with respect to the system size, $\xi, \ell = \mathcal{O}(1)$. Nevertheless, nearly all properties of phases of matter remain robust up to values just smaller than the system radius, $\xi, \ell = o(n^{1/d})$, where d is the spatial dimension. This follows because the light-cone of such circuits is much smaller than the system size [82]. For example, this guarantees that any long-range correlations and entanglement (or absence thereof) will be precisely retained. Emergent quasiparticle excitations above the ground state will also retain their properties, after coarse-graining over a distance ξ. When studying the complexity of recognizing phases of matter, a crucial question is how algorithms scale with the range of correlations in the state of interest. Hence, in our work, we treat ξ and ℓ as variable parameters and study the complexity of phase recognition as a function of ξ, ℓ.

2.2 Summary of main results

Our main result establishes that the quantum computational complexity of recognizing quantum phases of matter grows exponentially in the correlation range. For technical reasons, the circuit depth ℓ for generating the states we consider is polynomial in the correlation range ξ; hence, the complexity grows as a stretched exponential in the circuit depth ℓ. This rapid growth implies that the complexity is super-polynomial in n even for small values of $\xi = \omega(\log n)$ or $\ell = \mathrm{poly}(\log n)$.

Theorem 1 (Hardness of recognizing quantum phases of matter)**.** *For any discrete on-site symmetry group G, distinguishing any quantum phase of matter from the trivial phase requires quantum algorithms with computational time scaling exponentially in the correlation range ξ. The computational time becomes super-polynomial in the system size n whenever $\xi = \omega(\log n)$.*

This lower bound on the hardness of phase recognition is optimal, since rigorous algorithms exist with runtime growing exponentially in ξ [27, 29]. These algorithms succeed by learning the entire circuit description of a shallow quantum state. Hence, our result proves that recognizing the phase of matter of a shallow quantum state is comparably hard to performing full state tomography. We discuss the extension of Theorem 1 to mixed states in Section 2.6.

Before proceeding to our constructions, let us provide a high-level overview of our proof of Theorem 1. We consider the following phase recognition task: one is given access to many copies of a pure quantum state $|\psi\rangle$ drawn from either the trivial phase or a non-trivial phase of matter. The goal is to recognize the phase of matter of $|\psi\rangle$ using arbitrary efficient quantum circuits and measurements on all copies of the state. To prove that this task is computationally hard, we consider the complexity of recognizing a *random example* of each phase of matter [30]. That is, for both the trivial and any non-trivial phase, we generate a random quantum state in the phase by applying a *symmetric pseudorandom unitary* U to the fixed point state, $|\psi\rangle = U|\psi_0\rangle$. Crucially, we show that such unitaries can be formed in very low circuit depths. By definition, applying a symmetric PRU

to any state produces a state that is indistinguishable from Haar-random by any efficient quantum computation. Hence, neither the random trivial state nor the random non-trivial state can be distinguished from Haar-random, which implies that the two states also cannot be distinguished from one another. Therefore, recognizing the phase of matter is quantum computationally hard.

The states generated in our proof are the ground states of parent Hamiltonians $UH_0U^\dagger$, where H_0 is the Hamiltonian of the fixed-point state and U is a low-depth symmetric PRU. If the fixed-point Hamiltonian contains terms with a constant locality and geometric range, then the Hamiltonian $UH_0U^\dagger$ will feature a locality and geometric range growing linearly in ξ. Hence, to achieve super-polynomial hardness in n, we require that the Hamiltonian locality grows weakly in n, as $\omega(\log n)$. This contrasts with standard physical settings, in which the Hamiltonian is nearly always 2-local, even when the range of correlations ξ is large. Addressing the scaling of phase recognition with the correlation range ξ for constant-local Hamiltonians remains an essential open question.

In what follows, we first present the core technical results behind our proof by describing our definitions and constructions of low-depth symmetric PRUs. We then return to quantum phases of matter and provide additional discussions and analysis corroborating our result. We provide an extension to mixed state phases of matter at the end of the section.

2.3 Symmetric pseudorandom unitaries (PRUs)

Let us now introduce our construction of symmetric PRUs. Pseudorandom unitaries (PRUs) are efficient random unitary ensembles that are indistinguishable from the Haar ensemble by any efficient quantum algorithm [31, 34]. Recently, the existence of PRUs that are secure against any subexponential-time quantum algorithms has been proven [34] under a widely-accepted cryptographic conjecture: the quantum hardness of solving learning with errors (LWE) [88].

In the presence of a symmetry group, one can naturally consider the symmetric Haar ensemble, containing all unitaries that commute with the symmetry operation [70–76]. A symmetric PRU is then defined as an efficient ensemble of symmetric random unitaries that is indistinguishable from the symmetric Haar ensemble by any bounded-time quantum computation. The existence of symmetric PRUs does not follow from the existence of standard PRUs, and, to date, no constructions of symmetric PRUs have been established.

Our first main technical result is a construction of symmetric PRUs for any discrete on-site symmetry in $\mathrm{poly}(n)$ circuit depth. We will improve the circuit depth to $\mathrm{poly}(\log n)$ in the next subsection.

Theorem 2 (Symmetric pseudorandom unitaries)**.** *Under the conjecture that no subexponential-time quantum algorithm can solve LWE, symmetric pseudorandom unitaries with security against any* $\exp(o(n))$*-time quantum adversary exist, for every discrete on-site symmetry group* G*. The unitaries can be formed in polynomial circuit depth on any architecture. When* G *is Abelian, the unitaries can be compiled from individually symmetric* $\mathcal{O}(1)$*-local gates.*

Here, an on-site symmetry refers to any symmetry whose representation on n qudits is a tensor product, $R_g = \bigotimes_i R_g^i$, for all $g \in G$. We refer to Appendix A.1 for a detailed introduction to the representation theory of symmetry groups in many-body quantum systems. The proof of Theorem 2 is described below and given formally in Appendix B.

Construction overview. Our construction of symmetric PRUs for Theorem 2 proceeds in three steps. First, we decompose any n-qudit Hilbert space into a tensor sum of irreducible representations ("irreps") of the symmetry group, $\mathcal{H} = \bigoplus_\lambda (\mathcal{M}_\lambda \otimes \mathcal{F}_\lambda)$ [Fig. 2(a)]. Each $\mathcal{F}_\lambda$ is the Hilbert space of the irrep λ, and each $\mathcal{M}_\lambda$ captures the multiplicity of the irrep. In this basis, a symmetric Haar-random unitary is simply a tensor sum of random unitaries U_λ in each symmetry sector, $U = \bigoplus_\lambda (U_\lambda \otimes \mathbb{1}_{\mathcal{F}_\lambda})$.

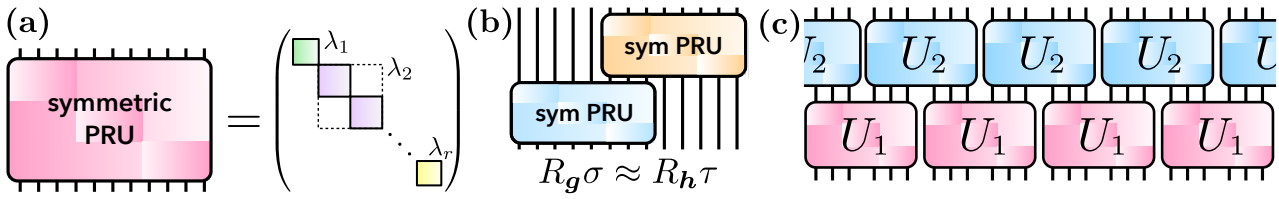

Figure 2: Schematic of our main results on symmetric pseudorandom unitaries (PRUs). **(a)** Any symmetric unitary is block-diagonal in the irreducible representations (irreps) $\lambda_1, \ldots, \lambda_r$ of the symmetry group. Our poly n-depth construction of symmetric PRUs applies independent random unitaries in each block, by performing a sequence of controlled PRUs conditioned on the irrep of the entire system (Theorem 2). **(b)** To construct symmetric PRUs in poly$(\log n)$ depth, we first prove that two small symmetric PRUs can "glue" into a symmetric PRU on a larger system (Lemma 1). This is used to prove that the two-layer circuit in which each unitary is drawn independently from a symmetric PRU ensemble is a low-depth symmetric PRU (Theorem 3). At a technical level, the symmetric gluing lemma follows by identifying the terms $R_{\boldsymbol{g}}\sigma$ and $R_{\boldsymbol{h}}\tau$ arising from the first and second unitaries' twirls, where $\boldsymbol{g}, \boldsymbol{h} \in G^{\otimes k}$ are symmetry group elements and $\sigma, \tau \in S_k$ are permutation operators. **(c)** A translation-invariant modification of the two-layer circuit, which underlies our construction of poly$(\log n)$-depth symmetric PRUs with translation invariance (Theorem 4).

Second, following this decomposition, we implement a sequence of *controlled pseudorandom unitaries* U_λ conditioned on the value of the system irrep λ. This directly yields U as above. Prior to our work, the existence of controlled PRUs had been widely believed but not formally proven in the literature. Hence, to implement this step, we prove that controlled PRUs exist under standard cryptographic assumptions and can be formed in polynomial circuit depth on any circuit geometry[2]. Our construction applies to arbitrary tensor product Hilbert spaces, which is necessary due to the variable dimensionality of $\mathcal{M}_\lambda$.

Third, we prove that the entire construction in steps one and two can be implemented using *individually symmetric* local gates whenever G is Abelian. Without this step, our construction would still yield symmetric PRUs in polynomial circuit depth. However, the individual gates in the circuit might not be symmetric (only the circuit as a whole would be symmetric). This is undesirable in several applications. To resolve this, we prove a general result: any Abelian symmetric unitary U can be implemented in polynomial depth via geometrically 5-local symmetric gates whenever the unitary of each irrep U_λ can be implemented in polynomial depth via geometrically 2-local gates. This result uses a mild connectivity assumption on the circuit architecture.

2.4 Symmetric PRUs in extremely low depth

We now present our second main technical result, which shows that symmetric PRUs can be realized in extremely low circuit depths.

Why low-depth symmetric PRUs may not exist. Before proceeding, we recall a potential counterargument against the possibility of low-depth symmetric random unitaries. Symmetries of physical systems lead to the existence of conserved charges, which label how the system responds to the symmetry operation. In a system with geometrically-local dynamics, charges are also conserved locally. Thus, one might question: Is it possible for a low-depth symmetric unitary to shuffle charges around the system in the same manner as a deep random unitary? This argument was used to

[2]During the preparation of this work, Tang and Wright [89] provided a distinct and independent proof of the existence of controlled pseudorandom unitaries.

rule out low-depth random unitaries in systems with any *continuous* symmetry [77]. In contrast, in what follows, we show that low-depth random unitaries can be naturally realized in systems with any discrete on-site symmetry. Intuitively, charges of discrete symmetries can be locally created and annihilated in neutral combinations, which allows the dynamics to mimic those of a deep random unitary even at very low depths.

At a high level, our construction of low-depth symmetric PRUs proceeds similarly to the construction of low-depth PRUs without symmetry [30]. We organize a system of n qudits [90] along a 1D line, and consider a two-layer brickwork circuit composed of symmetric PRUs acting on small patches of 2ξ qudits each. A translation-invariant version of this construction is shown in Fig. 2(c). Without translation invariance, our construction uses different small random unitaries throughout each circuit layer instead of the same small random unitary. We prove that the resulting two-layer circuit forms a symmetric PRU on all n qudits by proving that one can "glue" the small PRUs together one at a time. Namely, we prove the following gluing lemma for symmetric random unitaries [Fig. 2(b)]:

Lemma 1 (Gluing symmetric random unitaries)**.** *For any approximation error $\varepsilon \leq 1$, suppose each small unitary in the two-layer brickwork ensemble $\mathcal{E}$ is drawn from the symmetric Haar ensemble on 2ξ qudits. Then $\mathcal{E}$ forms an ε-approximate symmetric unitary k-design whenever $\xi \geq \log_2(nk^2|G|/\varepsilon)$.*

The lemma references the notion of a symmetric unitary design, which we review in Appendix A. By applying the gluing lemma n/ξ times [30], we find that the two-layer circuit forms a symmetric PRU whenever $\xi = \omega(\log n)$ (Appendix B). Moreover, as one increases ξ above this value, the security becomes exponentially stronger in the value of ξ.

Theorem 3 (Symmetric random unitaries in extremely low depth)**.** *Under the conjecture that no subexponential-time quantum algorithm can solve LWE, symmetric PRUs on n qudits with security against any $\exp(o(\xi))$-time adversary can be constructed with light-cone size $\mathcal{O}(\xi)$ and circuit depth $\mathrm{poly}(\xi)$, for any discrete on-site symmetry group G and $\xi = \omega(\log n)$.*

The circuit depth follows by instantiating each small PRU using Theorem 2, which yields a circuit depth $\ell = \mathrm{poly}(\xi)$. We have $\ell = \mathrm{poly}(\log n)$ if $\xi = \mathrm{poly}(\log n)$. As in Theorem 2, when G is Abelian, each unitary can be compiled from individually symmetric local gates.

Proof ideas for symmetric gluing lemma Our proof of the symmetric gluing lemma generalizes a large amount of technical machinery for random unitary analysis to the symmetric setting. Although seemingly straightforward, this generalization requires several careful choices to enable a tractable analysis. For the expert audience, we now summarize a few key intermediary steps and results. We begin by deriving an exact formula for the twirl of a symmetric Haar-random unitary in terms of a sum over permutation operators and tensor products of symmetry operators, $R_{\boldsymbol{g}} = \bigotimes_{i=1}^{k} R_{g_i}$. Then, we prove a much simpler *approximate* formula for the twirl:

$$\mathop{\mathbb{E}}_{U}\left[(U^{\otimes k})\, A\, (U^{\otimes k})^{\dagger}\right] \approx \frac{1}{D^k} \sum_{\sigma \in S_k} \sum_{\boldsymbol{g} \in G^{\otimes k}} \mathrm{tr}\big(A\sigma^{-1} R_{\boldsymbol{g}}^{-1}\big) \cdot R_{\boldsymbol{g}}\sigma. \tag{2.1}$$

This replaces the sum over permutation operators in the approximate Haar twirl without symmetry [30] with a sum over both permutation and symmetry group operators. Third, we prove a key technical lemma for translating between different notions of error in symmetric unitary designs. Finally, from these results, we provide a self-contained proof that ε-approximate symmetric unitary k-designs can be realized in $\mathrm{poly}(k) \cdot \mathrm{poly}\log(n/\varepsilon)$ depth for any discrete symmetry group G, without relying on any spectral gap analysis. We expect that each of these results may be useful in future works, and refer the interested reader to Appendix A for a full exposition and further details.

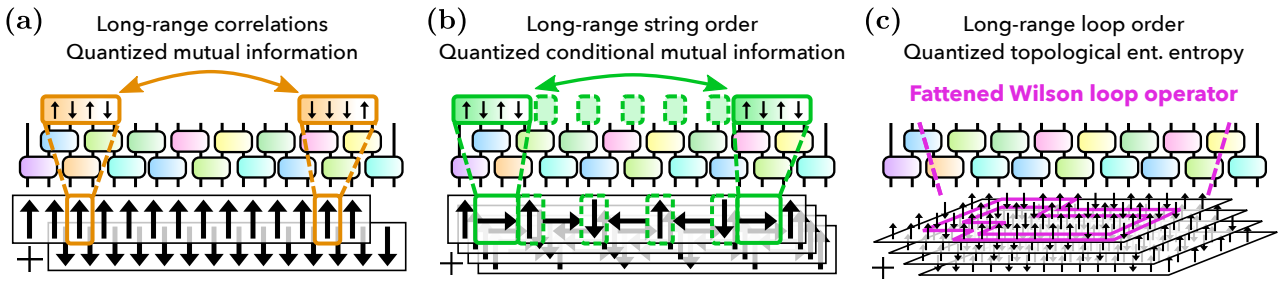

Figure 3: Three phases of matter and their order parameters after application of a low-depth symmetric PRU. **(a)** A symmetry-breaking phase (example: a GHZ state) features long-range correlations and a quantized non-zero mutual information. **(b)** A symmetry-protected topological (SPT) phase (example: a 1D cluster state) features string order parameters and a quantized non-zero conditional mutual information. **(c)** Topological orders (example: a toric code state) feature loop order parameters and a quantized non-zero topological entanglement entropy. In all three cases, the order parameters and quantized correlations are retained after the low-depth PRU. However, they cannot be detected by any polynomial-time observer without knowledge of the unitary that was applied.

Translation-invariant symmetric PRUs. Our third main technical result establishes the existence of *translation-invariant* symmetric PRUs. Beyond on-site symmetry groups, translation symmetry is by far the most common symmetry in the physical world and plays a crucial simplifying role in many quantum phenomena. Our translation-invariant random unitary ensemble corresponds to the same two-layer circuit, but in which each unitary in the first layer is identical, and similarly for the second layer [Fig. 2(c)]. This yields a symmetric random unitary that is invariant under translations by $a = 2\xi$ lattice sites. We prove that the resulting unitary is indistinguishable from a translation-invariant symmetric Haar-random unitary by any efficient quantum algorithm.

Theorem 4 (Translation-invariant symmetric pseudorandom unitaries)**.** *Under the conjecture that no subexponential-time quantum algorithm can solve LWE, one-dimensional translation-invariant symmetric PRUs with security against any* $\exp(o(\xi))$*-time adversary can be constructed with light-cone size* $\mathcal{O}(\xi)$ *and circuit depth* $\mathrm{poly}(\xi)$*, for any discrete on-site symmetry group* G *and* $a = 2\xi = \omega(\log n)$*.*

As in Theorems 2 and 3, when G is Abelian, the translation-invariant symmetric PRUs can be compiled from individually symmetric local gates. Extending this result to translations over a constant distance, $a = \mathcal{O}(1)$, remains an interesting open direction. The two-layer circuit explicitly breaks translation symmetry for any $a < 2\xi$.

Our proof of Theorem 4 is provided in Appendix D and proceeds along a similar path to Lemma 1. The primary difference is that the proof involves a significantly more intricate bound over a huge sum of permutation operators. We introduce several helpful methods for counting permutations that enable us to complete this bound.

2.5 Hardness of recognizing pure state quantum phases of matter

Having established the existence of low-depth symmetric PRUs, we now return to their implications for recognizing quantum phases of matter. Our main hardness result, Theorem 1, follows directly from combining our symmetric PRU constructions (Theorems 3 and 4) with the proof strategy outlined in Section 2.2. The key insight is that symmetric PRUs allow us to generate quantum states in well-defined phases that are computationally indistinguishable from Haar-random states, making phase recognition computationally hard. Complete proof details are provided in Appendix E.

In the remainder of this section, we explore the implications of this result through concrete examples, provide intuition for why phase recognition becomes hard, and analyze how our results apply to various extensions and practical scenarios.

Concrete examples. To illustrate the power of our hardness result, consider the following four quantum states on n qubits with symmetry $G = \mathbb{Z}_2 \times \mathbb{Z}_2$. Here, U denotes a low-depth symmetric PRU from our construction. Our results prove that the four states, each representing a different phase of matter, are all computationally indistinguishable from one another:

1. $|\psi_{\text{triv}}\rangle = U\,|\text{product}\rangle$, where $|\text{product}\rangle$ is a product state respecting both symmetries. This state is in the trivial phase.

2. $|\psi_{\text{SSB}}\rangle = U\,|\text{GHZ}\rangle$, where $|\text{GHZ}\rangle$ is a GHZ state that spontaneously breaks either one or both $\mathbb{Z}_2$ symmetries. This state is in a spontaneous symmetry-broken phase [4].

3. $|\psi_{\text{SPT}}\rangle = U\,|\text{cluster}\rangle$, where $|\text{cluster}\rangle$ is the 1D cluster state [4]. This state is in a symmetry-protected topological phase.

4. $|\psi_{\text{topo}}\rangle = U\,|\text{TC}\rangle$, where $|\text{TC}\rangle$ is the toric code state and respects both symmetries. This state is in a topological phase [4].

Each state is in a well-defined and distinct phase of matter from all other states. The different phases of matter are each characterized by their own unique patterns of entanglement and long-range correlations. If an observer has knowledge of the unitary U, then the observer can easily use this knowledge of U to detect these correlations, and verify that each state is in its given phase. However, our results prove that without knowledge of U (i.e. when the quantum state is unknown to the observer), no efficient quantum algorithm can distinguish between any of the four families of states.

Why does phase recognition become hard? While our theorem formally establishes that recognizing phases of matter is hard, our technical proof does not provide much intuition as to why this is the case. Let us address this intuition through the lens of one common approach to phase recognition: measuring order parameters (Fig. 3). Order parameters are observables whose expectation values diagnose a particular phase of matter. For example, symmetry-breaking phases feature long-range correlations with respect to an order parameter that breaks the symmetry. For the GHZ state, these correspond to the expectation values, $\langle\psi|\, Z_j Z_i\, |\psi\rangle$, for distant sites i, j. Similarly, one-dimensional SPT phases feature string-like order parameters, which are equal to a symmetry operator in their bulk while breaking a symmetry at each edge. For the cluster state, these correspond to the expectation values, $\langle\psi|\, Z_i X_{i+1} X_{i+3} \dots X_{j-3} X_{j-1} Z_j\, |\psi\rangle$, for $j = i \bmod 2$. While the order parameters of many toy-model systems happen to be simple, this is not guaranteed to be the case.

The key intuition behind our results is that, without prior knowledge about a quantum state, its order parameters can be quantum computationally hard to identify. For example, when considering spontaneous breaking of a $\mathbb{Z}_2$ symmetry, there is no property of the symmetry or phase of matter that fixes the order parameter to be a simple one-body Z operator (as in the GHZ state). Rather, after conjugation by a low-depth symmetric PRU, all possible order parameters become scrambled superpositions of many symmetry-breaking operators acting on regions of $\mathcal{O}(\xi)$ qubits [Fig. 3(a)]. The overlap of these scrambled order parameters with any simple one-body operator is typically exponentially small in ξ, which makes the phase of matter hard to recognize using simple one-body expectation values. At the same time, the full $\mathcal{O}(\xi)$-body order parameters contain exponentially many random and non-commuting terms, and hence are hard to reconstruct in any experiment without knowledge of U. This causes order-parameter-based approaches for recognizing phases of matter to

require exponential resources as ξ increases. Analogous arguments apply to the order parameters of trivial, SPT, and topological phases; see Fig. 3(b,c) for an illustration.

What about entanglement-based approaches? Beyond order parameters, phase recognition often relies on measuring characteristic entanglement signatures that distinguish different phases of matter. However, these approaches encounter the same fundamental scaling limitations with the correlation range ξ. Consider the Kitaev-Preskill construction [91] for measuring the topological entanglement entropy (TEE): $S_{\text{topo}} = S_{AB} + S_{BC} + S_{AC} - S_A - S_B - S_C - S_{ABC}$. This quantity is a powerful tool designed to distinguish topological phases ($S_{\text{topo}} > 0$) from trivial phases ($S_{\text{topo}} = 0$); see [92] for a rigorous version of TEE that addresses spurious cases [93, 94]. When the correlation range is ξ, the topological entanglement entropy needs to be measured on regions of radius $\mathcal{O}(\xi)$. The sample complexity for measuring entropies on such regions grows exponentially in ξ. Similarly, the mutual information $I(A : B) = S_A + S_B - S_{AB}$ for detecting long-range correlations in symmetry-breaking phases, and the conditional mutual information $I(A : B|C)$ for capturing long-range string correlations in SPT phases, both require measuring entropies on regions with radius $\mathcal{O}(\xi)$. This leads to the same exponential scaling in the sample complexity, confirming that entanglement-based phase recognition approaches suffer from the same fundamental hardness as order-parameter approaches.

Implications for existing phase recognition algorithms. Our hardness results establish the optimality of existing rigorous algorithms for phase recognition [27–29], which achieve runtime exponential in ξ. This proves that no algorithm can fundamentally improve upon this scaling. On the other hand, most common approaches to phase recognition are highly heuristic and leverage assumptions and prior knowledge about the quantum state of interest [17–20, 22]. Our results show that there exist quantum states where all such approaches must fail; we emphasize that this does not preclude their success in specific practical settings.

A recent class of algorithms based on the renormalization group represents a particularly interesting case [23–25], as these methods were specifically designed to overcome scaling limitations. These algorithms operate in a coarse-grained manner, first observing local properties and gradually zooming out to larger scales to determine the phase of matter. Building on physical assumptions, these algorithms were shown to learn the entire circuit descriptions of quantum ground states in time scaling only logarithmically in the correlation length. Our results show that this logarithmic scaling cannot hold in complete generality, as our Theorem 1 establishes that phase recognition fundamentally requires exponential scaling. The reason these renormalization-group-based algorithms must fail in general is that there is no information to be gleaned from local properties of locally-scrambled random quantum states. This follows because locally-scrambled quantum states have local density matrices that concentrate to the maximally mixed state with errors decaying exponentially in ξ and the circuit depth ℓ [30, 69]. Without specific structures or physical restrictions, phase recognition is fundamentally hard and cannot achieve the logarithmic scaling claimed by these approaches.

What about non-symmetric ground states? In our proof of Theorem 1 and the examples above, we focus on states $|\psi\rangle$ that are symmetric in each phase. This is always possible because there exists a symmetric representative of any phase of matter, since the ground state of a symmetric Hamiltonian is symmetric. However, our results also imply strong limitations on recognizing non-symmetric representatives of a phase as well. For example, consider the state $U|0^n\rangle$, which explicitly breaks the $\mathbb{Z}_2$ symmetry $\prod_i X_i$, instead of $U|\text{GHZ}\rangle$ where $|\text{GHZ}\rangle = (|0^n\rangle + |1^n\rangle)/\sqrt{2}$, which preserves the symmetry. Our results prove that the only information one can efficiently extract from the state is the total population in each symmetry sector; see Appendix E for details. This allows distinguishing explicit symmetry-breaking states like $U|0^n\rangle$ from the trivial phase. However, it still

rules out distinguishing such states from non-symmetric states in other phases of matter, such as SPT phases [95]. Moreover, Theorem 1 implies that *obtaining* the explicitly symmetry-broken state $U|0^n\rangle$ from the symmetric state $U|\mathrm{GHZ}\rangle$ is computationally hard without knowledge of U. If this were not true, then one could efficiently learn the phase of matter.

What about constant-local translation-invariant Hamiltonians? The most restrictive definitions of quantum phases of matter focus on an extremely small subset of states: the ground states of constant-local translation-invariant Hamiltonians with a constant spectral gap. One might wonder, what is the complexity of recognizing phases of matter in such states? We provide a short argument below that, from a complexity viewpoint, phases of matter become trivial to recognize in this restricted setting. This follows because there are only polynomially many possible ground states of such Hamiltonians, which enables a simple polynomial-time brute force search algorithm. Hence, to formulate phase recognition as a non-trivial computational problem, one must expand the definition to include a larger-than-polynomial number of possible ground states.

The argument proceeds as follows. Suppose we would like to distinguish whether a quantum state is related to a given fixed point state by a gapped adiabatic path of k-local Hamiltonians that are translation-invariant over distance a, where $k, a = \mathcal{O}(1)$. Let Δ denote the minimum gap along the path; the inverse gap Δ^{-1} plays a role similar to ξ in our analysis. A k-local Hamiltonian is specified by at most $4^k = \mathcal{O}(1)$ continuous parameters. The constant spectral gap implies that one can discretize the parameter space by increments of Δ/n; the ground state does not change appreciably for smaller parameter changes. Hence, there are at most $(n/\Delta)^{\mathcal{O}(1)} = \mathrm{poly}(n, \Delta^{-1})$ possible ground states to consider. To measure the fidelity of the unknown quantum state with any one of these polynomial number of candidate ground states, one can implement an efficient quantum circuit that maps from the ground state to the trivial fixed point state, e.g. by Trotterizing an adiabatic time evolution, and then measure the fidelity with the fixed point state. A brute force search over all possible ground states solves the task in time $\mathrm{poly}(n, \Delta^{-1})$.

What about states generated from individually symmetric local gates? Following our discussion of symmetric PRUs, our results for Abelian symmetry groups are slightly stronger than those for non-Abelian groups. For Abelian symmetries, our results apply even when each individual constant-local gate in the unitary U is required to be symmetric. For non-Abelian symmetries, our results apply only when each individual 2ξ-local patch is symmetric [see Fig. 2(c)], but not necessarily each individual constant-local gate comprising each patch. To achieve this in the Abelian case, we require a mild assumption that one can perform symmetric gates between nearest-neighbor triplets of qudits on which the symmetry acts; see Appendix B. This avoids unusual counter-examples where no symmetric qudits can interact directly with one another [96].

2.6 Hardness of recognizing mixed state quantum phases of matter

We conclude this section by discussing the extension of our results to mixed state quantum phases of matter. Traditionally, mixed state phases have been studied in the context of finite-temperature Gibbs states of quantum or classical Hamiltonians [1]. They have also received renewed interest recently [97–105], sparked by experimental realizations affected by noise and decoherence.

The precise definition and classification of mixed state phases of matter remains an active research topic [97–105]. One common definition proceeds as follows [100–103]. Two mixed quantum states ρ_1 and ρ_2 are in the same phase of matter if and only if there exists a pair of local channels, $\mathcal{C}_1$ and $\mathcal{C}_2$, such that $\mathcal{C}_1(\rho_1) = \rho_2$ and $\mathcal{C}_2(\rho_2) = \rho_1$. The use of quantum channels instead of unitary transformations is needed to enact equivalence transformations between mixed states with different

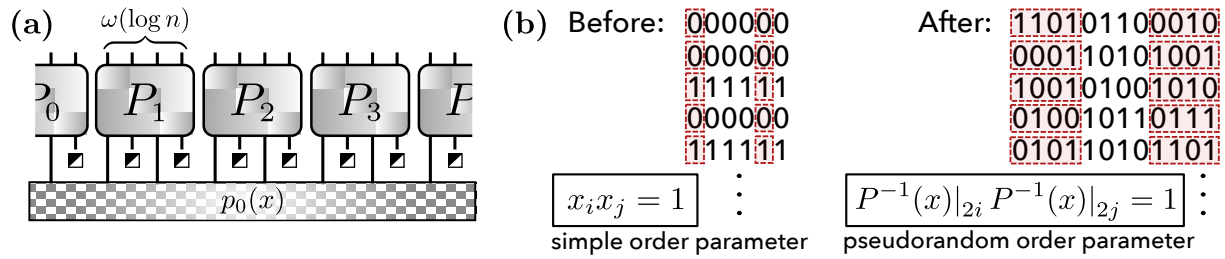

Figure 4: Extension of our results to classical phases of matter. **(a)** We consider a reversible one-layer brickwork circuit composed of symmetric pseudorandom permutations acting on patches of $\xi = \omega(\log n)$ bits each. The input to the circuit is any fixed point probability distribution, $p_0(x)$, tensored with n mixed ancilla bits (black-white squares). We prove that for any input $p_0(x)$, the output of the circuit is indistinguishable from the maximally mixed distribution, i.e. the trivial phase. **(b)** Samples from the $\mathbb{Z}_2$-symmetry-breaking fixed point distribution before and after the circuit. The circuit transforms the order parameter from a simple single-body observable into a pseudorandom function on $2\xi = \text{poly}(\log n)$ bits. This hides the phase of matter from any poly n-time observer.

spectra. The use of two channels instead of one follows because quantum channels are, in general, not invertible. If the system is symmetric, then one also requires that each channel is covariant under the symmetry transformation, $\mathcal{C}(\cdot) = R_g^\dagger \mathcal{C}(R_g(\cdot)R_g^\dagger)R_g$ for all $g \in G$. It was recently argued that additional smoothness constraints on the channels $\mathcal{C}_1$ and $\mathcal{C}_2$ should also be imposed [105].

Similar to the pure state setting, our results apply to any definition of mixed state phases of matter such that the following holds. For each phase of matter, we can associate a representative fixed point mixed state ρ_0. Then, we assume that any state ρ that is related to ρ_0 by symmetric geometrically-local quantum channels as above is in the same phase as ρ_0. We let ℓ denote the maximum allowed circuit depth of the channels and ξ the light-cone radius. We require only that the allowed channels include symmetric geometrically-local unitary circuits, and tensor product operations with ancilla qubits that are maximally mixed and unentangled with the system.

Our main result shows that mixed state phases of matter are as hard to recognize as pure state phases of matter. Namely, we prove that it is quantum computationally hard to distinguish any mixed state phase of matter from a mixture of maximally mixed states in each symmetry charge sector. The latter should be viewed as belonging to the trivial phase. The population in each charge sector is inherited from the fixed point of the mixed state phase.

Theorem 5 (Hardness of recognizing mixed state phases of matter)**.** *For any discrete on-site symmetry group G, distinguishing any mixed state phase of matter from a mixture of maximally mixed states in each symmetry charge sector requires quantum computational time scaling exponentially in the correlation range ξ. The time becomes super-polynomial in the system size n whenever $\xi = \omega(\log n)$.*

This theorem holds for all types of mixed state phases of matter, including both topological and symmetry-breaking mixed state phases. This complements recent results on the hardness of recognizing strong-to-weak spontaneous symmetry breaking [106]. As with Theorem 1, our proof of Theorem 5 follows directly from our construction of low-depth symmetric PRUs.

3 Classical phases of matter

While our primary focus is on quantum phases of matter, we find that, surprisingly, the hardness of recognizing phases of matter extends to purely classical states and phases as well. The hallmark

example of a classical phase of matter occurs in the two-dimensional Ising model, $H = \sum_{\langle i,j \rangle} Z_i Z_j$, at finite temperature $1/\beta$ [1]. The Gibbs state defines a classical probability distribution,

$$p_\beta(x) = \frac{\langle x| \, e^{-\beta H} \, |x\rangle}{\sum_x \langle x| \, e^{-\beta H} \, |x\rangle}, \tag{3.1}$$

where $x \in \{0,1\}^n$. At high temperatures ($\beta < \beta_c \approx 0.441$), the distribution is in the trivial phase and features only short-range correlations. At low temperatures ($\beta > \beta_c$), the distribution becomes symmetry-breaking, leading to long-range correlations across the entire system.

Similar to quantum systems, we can formalize the definition of classical phases of matter through local circuit equivalences. Following our discussion of mixed state quantum phases of matter, two classical probability distributions $p_1(x)$ and $p_2(x)$ are in the same phase of matter if and only if there exists a pair of geometrically-local stochastic maps $\mathcal{M}_1$ and $\mathcal{M}_2$ such that $\mathcal{M}_1(p_1) = p_2$ and $\mathcal{M}_2(p_2) = p_1$. We let ℓ denote the circuit depth of the maps and ξ the light-cone radius. This precisely mirrors the definition of mixed state quantum phases of matter, but with all objects restricted to be purely classical. For systems possessing a symmetry, both maps must be covariant under the symmetry operation. We refer to Appendix F for further details.

Our main result establishes the computational hardness of recognizing classical phases of matter:

Theorem 6 (Hardness of recognizing classical phases of matter)**.** *For any discrete on-site symmetry group G, recognizing whether a classical probability distribution is in a trivial versus symmetry-breaking phase requires a quantum or classical computational time scaling exponentially in the correlation range ξ. The time becomes super-polynomial in the system size n whenever $\xi = \omega(\log n)$.*

We focus on symmetry-breaking phases since topological phases are not present in classical systems.

As in the quantum setting, we prove Theorem 6 by considering the complexity of recognizing a random example of each classical phase of matter. Our random example is generated by applying a random reversible circuit $\mathcal{M}$ to a fixed point distribution $p_0(x)$ of each phase, as depicted in Fig. 4(a). The circuit $\mathcal{M}$ consists of only a single layer: a tensor product of $2n/\xi$ symmetric pseudorandom permutations (PRPs) [107] applied to disjoint patches of ξ bits across the system. Before applying the circuit $\mathcal{M}$, we tensor the fixed point distribution $p_0(x)$ with n ancilla bits (one per physical bit), each initialized to be uniformly random (i.e. in the maximally mixed state, $\mathbb{1}/2$). This operation adds entropy to the system without changing its phase of matter.

We prove Theorem 6 by showing that for any fixed point distribution $p_0(x)$, the resulting random ensemble is exponentially hard to distinguish from the maximally mixed distribution by any classical or quantum observer. The intuition behind our result parallels the quantum case; nonetheless, it is somewhat unexpected given that classical circuits in one-dimension cannot appear globally random in any depth sublinear in the system size [30]. Consider the fixed point distribution for the $\mathbb{Z}_2$-symmetry-broken phase: $p_0(x) = \frac{1}{2}\delta_{x,0^n} + \frac{1}{2}\delta_{x,1^n}$. This features long-range correlations $\langle f_i(x) f_j(x) \rangle = 1$ with respect to the order parameter $f_i(x) = (-1)^{x_i}$. After conjugating by PRPs on each patch, the distribution still features long-range correlations, but now with respect to a *scrambled* order parameter $f'_i(x) = f_i(P^{-1}(x))$, where P is the permutation acting on the patch containing bit i. Since P is a PRP, the scrambled order parameter f'_i is a pseudorandom function on the patch, making it very difficult to learn for any classical or quantum algorithm. Hence, as in the quantum case, the random probability distribution maintains genuine long-range symmetry-breaking order, yet the parameter signaling this order is impossible to detect efficiently.

The technical details of our proof are relatively straightforward. We consider the distribution of bitstrings $x^{(1)}, x^{(2)}, \ldots, x^{(k)}$ obtained by querying the unknown distribution $k = \text{poly}(n)$ times. Due to the entropy of the ancilla bits, the k bitstrings are *distinct* from one another, i.e. $x^{(i)} \neq x^{(j)}$ for all $i < j$ [32], with high probability. Moreover, this distinctness also holds at the level of each

individual patch: the set of k ξ-bit strings obtained by restricting $x^{(1)}, x^{(2)}, \ldots, x^{(k)}$ to any patch of ξ bits is distinct with high probability whenever $\xi = \omega(\log n)$ [44]. The key observation is that the probability of any given set of k patchwise-distinct bitstrings appearing is uniform across all such sets due to the pseudorandom permutations on each patch. Therefore, the distribution of interest cannot be distinguished from the uniform distribution over all sets of k bitstrings, which corresponds to the trivial phase of matter. We refer to Appendix F for complete details.

Our classical construction relies crucially on the probability distributions of interest having a high local entropy, which is provided by the maximally mixed ancilla bits. This high local entropy allows one to locally hide information about the order parameters of the phase. This intuition suggests that a similar one-layer tensor-product circuit should be able to hide the phase of matter of highly-mixed quantum states as well. However, this approach cannot extend to pure quantum states, where a simple purity test would distinguish any state generated by any tensor-product circuit from a truly random state. One could also measure mutual information between distant patches of the quantum state to enable phase recognition. Hence, the hardness of recognizing *pure state* quantum phases of matter arises crucially from the ability of quantum circuits to appear globally random at very low depth, which does not occur in classical systems.

4 Discussions

Our results raise several questions regarding both the recognition of phases of matter, and symmetric random unitaries more broadly. First, as discussed in the introduction, how does one reconcile our results with the observation that phases of matter are often efficiently recognizable in practice? Is there a fundamental property of the ground states encountered in the physical world that guarantees this efficiency? In particular, the most natural limitation in our results is that the locality of the parent Hamiltonian grows with the correlation range ξ. However, even a 2-local Hamiltonian can possess ground states with arbitrarily large correlation ranges ξ, since the ξ will typically scale with the inverse of the spectral gap. This raises a crucial question: Can one prove that the ground states of constant-local Hamiltonians exhibit different physical behaviors that enable efficient phase recognition with computational time scaling only polynomially in the correlation range ξ? And if so, does this hold for every ground state, or only for a large class typically encountered in practice?

Indeed, in the same direction, arguments from conformal field theory suggest that the complexity of recognizing many phases of matter should grow only polynomially in the correlation length, and that this scaling can be achieved by measuring extremely simple fixed point order parameters [108]. For example, the magnetization of a spin chain (a one-body Pauli operator) is predicted to decay polynomially in the correlation length ξ as one approaches a phase transition. Thus, one can detect the phase by measuring the magnetization with only a polynomial number of samples. This prediction is tremendously simplifying, yet it goes far beyond our current mathematical understanding of quantum ground states and phases of matter via short-time adiabatic evolutions. Is it possible to improve our understanding to encompass these simplifying properties?

Finally, beyond phases of matter, our results make progress on extending ideas from pseudorandom unitaries to physical systems that appear in the natural world. This provides evidence that other important properties of random unitaries, such as quantum sampling advantages [51, 52], barren plateaus in variational algorithms [109], and classical shadow tomography [110], should also persist at low depths in quantum systems with any discrete symmetry. An important remaining question concerns quantum systems with continuous symmetries, such as charge or energy conservation. While simple light-cone arguments rule out the complete formation of low-depth random unitaries in these settings [77], is it possible that other features of low-depth random unitaries might persist?

Acknowledgments:

We are grateful to Shankar Balasubramanian, Soonwon Choi, Samuel J. Garratt, Soumik Ghosh, Jeongwan Haah, Jonas Haferkamp, Isaac Kim, Ethan Lake, Yu-Jie Liu, John Preskill, Norbert Schuch, and Nathanan Tantivasadakarn for valuable discussions and insights. T.S. acknowledges support from the Walter Burke Institute for Theoretical Physics at Caltech. T.S. and H.H. acknowledge support the U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers, Quantum Systems Accelerator. The Institute for Quantum Information and Matter, with which T.S. and H.H. are affiliated, is an NSF Physics Frontiers Center (NSF Grant PHY-2317110). H.H. acknowledge support from the Broadcom Innovation Fund.

Appendices

Our Appendices are organized as follows. In Appendix A, we present our main results on symmetric Haar-random unitaries and designs. In Appendix B, we present our main results on symmetric pseudorandom unitaries. In Appendix C, we present our construction of controlled pseudorandom unitaries, which is used in our construction of symmetric pseudorandom unitaries in Appendix B. In Appendix D, we present our results on translation-invariant random unitaries. In Appendix E, we present the implications of these results for recognizing quantum phases of matter. In Appendix F, we present our independent results on the hardness of recognizing classical phases of matter.

Contents

A Symmetric random unitaries **16**
A.1 Background on representations of symmetry groups 16
A.2 The exact symmetric Haar twirl 18
A.3 Approximate symmetric Haar twirl 21
A.4 Symmetric unitary designs 24
A.5 Proof of Lemma 1: Gluing symmetric random unitaries 24

B Symmetric pseudorandom unitaries **26**
B.1 Definition of symmetric PRUs 26
B.2 Proof of Theorem 2: Symmetric PRUs in polynomial depth 27
B.3 Proof of Theorem 3: Symmetric PRUs in poly-logarithmic depth 32

C Controlled pseudorandom unitaries **32**
C.1 Summary of results 32
C.2 Proof of Theorem 10: Controlled-PFC ensemble 34
C.3 Proof of Theorem 11: Controlled-LRFC ensemble 37

D Translation-invariant random unitaries **40**
D.1 Approximate translation-invariant Haar twirl 40
D.2 Translation-invariant PRUs in extremely low depth 41
D.3 Symmetric translation-invariant PRUs in extremely low depth 44

E Quantum phases of matter **46**
E.1 Proof of Theorem 1: Hardness of recognizing pure state phases of matter 46

E.2 Proof of Theorem 5: Hardness of recognizing mixed state phases of matter 47

F Classical phases of matter **48**
F.1 Proof of Theorem 6: Hardness of recognizing classical phases of matter 48

A Symmetric random unitaries

In this section, we present the full details of our results on gluing symmetric random unitaries. We begin in Appendix A.1 with a basic introduction to symmetry groups and their representations on qudit Hilbert spaces. We then turn to random symmetric unitaries in Appendix A.2, and provide an explicit formula for the twirl over a Haar-random symmetric unitary. In Appendix A.3, we present our first main result by providing a much simpler approximate expression for the symmetric Haar twirl. Both our exact and approximate expressions hold for any unitary symmetry group, discrete or continuous. In Appendix A.4, we provide the definition of a symmetric unitary k-design, following the standard definition for unitary k-designs without symmetry. In Appendix A.5, we present our second main result by proving that random symmetric unitaries "glue" together whenever they overlap on a small region. Our gluing lemma applies to any discrete on-site symmetry group, and forms the basis of our results on generating low-depth symmetric pseudorandom unitaries in Appendix B.

A.1 Background on representations of symmetry groups

We begin by providing a short introduction to the representation theory of symmetry groups in quantum many-body systems. We refer to standard textbooks for a comprehensive review [111].

Unitary representations of symmetry groups.—Consider any group G and any finite Hilbert space $\mathcal{H}$. A unitary representation R of G on $\mathcal{H}$ is a homomorphism mapping each element $g \in G$ to a unitary operator R_g acting on $\mathcal{H}$. A standard fact of mathematics states that any unitary representation on a finite Hilbert space can be decomposed as a tensor sum of *irreducible representations*, or "irreps", as

$$\mathcal{H} = \bigoplus_\lambda \left(\mathcal{M}_\lambda \otimes \mathcal{F}_\lambda\right), \qquad R_g = \bigoplus_\lambda \left(\mathbb{1} \otimes R_g^\lambda\right). \tag{A.1}$$

Here, the tensor sum is over all irreps λ of the symmetry group. Each irrep λ corresponds to a representation R^λ of the group acting on a Hilbert space $\mathcal{F}_\lambda$. By definition, the irreducible representation cannot be decomposed into a tensor sum of any sub-representations. We denote the dimension of each irrep as $d_\lambda = |\mathcal{F}_\lambda|$. Each irrep might also appear many times in the decomposition. To capture this, each irrep is also associated with a Hilbert space $\mathcal{M}_\lambda$, whose dimension is equal to the number of times that the irrep appears, $D_\lambda \equiv |\mathcal{M}_\lambda|$.

In our work, we are interested in random symmetric unitary operators U on $\mathcal{H}$, which commute with all elements of the symmetry group. From Schur's lemma, any such unitary can be written in the block-diagonal form,

$$U = \bigoplus_\lambda \left(U_\lambda \otimes \mathbb{1}_{\mathcal{F}_\lambda}\right), \tag{A.2}$$

where each U_λ acts on $\mathcal{M}_\lambda$. The unitary acts trivially on the $\mathcal{F}_\lambda$ registers, which is necessary in order to commute with all symmetry group elements. One can also easily see that any unitary of the form above is symmetric, i.e. it commutes with all symmetry group elements.

Discrete symmetry groups.—In quantum many-body systems, the Hilbert space factorizes into a tensor product of small components, $\mathcal{H} = \bigotimes_{i=1}^n \mathcal{H}_i$, where each $\mathcal{H}_i$ is the Hilbert space of an individual qubit or qudit of the system. The most common symmetry operations in many-body physical systems fall into two categories. In the first, *on-site* symmetries, symmetry group elements act in a parallel on

all qudits at a time. That is, the representation, R_g, of a symmetry group element $g \in G$ on $\mathcal{H}$ is given by the tensor product of local representations of g on each local Hilbert space, $R_g = \bigotimes_{i=1}^n R_g^i$. In the second, *lattice* symmetries, symmetry group elements act to permute the qudits (and their associated Hilbert spaces) amongst one another. In this work, we will primarily focus on on-site symmetries. We extend our results to one particular lattice symmetry, translation-invariance, in Appendix D.

Let us now turn to the representation theory of discrete on-site symmetry groups in particular. Any finite discrete group possesses a finite number of irreducible representations. In our work, we focus for simplicity on the most common representation of a discrete symmetry group, known as the *regular* representation. As discussed in the following paragraphs, the tensor product of any representation of a symmetry group with a qudit in the regular representation yields a a larger regular representation. Hence, for our results on phases of matter, we can assume without loss of generality that the symmetry acts in the regular representation on each qudit. If it does not, then we can take the tensor product of each qudit with a symmetric ancilla qudit where the symmetry acts in the regular representation on the ancilla. This does not change the phase of matter, and yields a larger on-site qudit Hilbert space where the symmetry acts in the regular representation.

Let us now describe the regular representation and its properties. For any finite discrete group G, the regular representation R^{reg} of G acts on a Hilbert space, $\mathcal{F}_{\text{reg}} = \{|g\rangle \, | g \in G\}$. The Hilbert space has dimension $|G|$, and is spanned by computational basis states labeled by each element $g \in G$. The regular representation of G acts on these basis states via right multiplication, $R_g^{\text{reg}} |h\rangle = |gh\rangle$. One useful property of the regular representation is that the symmetry group operators, R_g^{reg}, are orthogonal to one another. In particular, we have

$$\frac{1}{|G|} \operatorname{tr}\Big(R_g^{\text{reg}} (R_h^{\text{reg}})^\dagger\Big) = \frac{1}{|G|} \operatorname{tr}\Big(R_{gh^{-1}}^{\text{reg}}\Big) = \delta_{g,h}, \tag{A.3}$$

for any $g, h \in G$. From this property, one can show that the regular representation has an irrep decomposition [Eq. (A.1)] in which every irrep λ appears with multiplicity equal to its dimension d_λ [111]. This fact also yields the convenient formula, $\sum_\lambda d_\lambda^2 = |G|$ [111].

As previously mentioned, a second useful property of the regular representation concerns its behavior under tensor products. Namely, the regular representation is a *fixed point* under the tensor product. To elaborate, consider the tensor product of the regular representation, $\mathcal{F}_{\text{reg}}, R_g^{\text{reg}}$, of a group G, with any other representation, $\mathcal{F}_2, R_g^2$, of G. From these two original representations, one can naturally define a third representation, $R_g = R_g^{\text{reg}} \otimes R_g^2$, which acts on the combined Hilbert space $\mathcal{F}_{\text{reg}} \otimes \mathcal{F}_2$. As for any representation, the tensor product representation can be decomposed as a sum of irreps. For a general pair of original representations, this decomposition can be complicated to determine. However, when either one of the input representations is regular, one finds that the tensor product representation has an irrep decomposition identical to that of the regular representation, up to an overall constant factor increase in the multiplicity. That is, the irrep decomposition of the tensor product representation contains each irrep λ of G with multiplicity $d_\lambda |\mathcal{F}_2|$. This implies that there exists an isomorphism, $\mathcal{F}_{\text{reg}} \otimes \mathcal{F}_2 \cong \mathbb{C}^{|\mathcal{F}_2|} \otimes \mathcal{F}_{\text{reg}}$, such that the symmetry group acts only on the rightmost register, $R_g = \mathbb{1}_{|\mathcal{F}_2|} \otimes R_g^{\text{reg}}$.

We can now return to the quantum many-body setting. Suppose that each local Hilbert space $\mathcal{H}_i$ transforms under the regular representation of G. That is, we can write $\mathcal{H}_i \cong \mathcal{F}_{\text{reg},i} \otimes \mathcal{A}_i$, where $R_g^i = R_g^{\text{reg},i} \otimes \mathbb{1}_{\mathcal{A}_i}$ and $\mathcal{A}_i$ is an ancilla register of arbitrary dimension d'. Repeating the tensor product procedure in the previous paragraph $n-1$ times, one finds that the many-body Hilbert space $\mathcal{H} = \bigotimes_{i=1}^n \mathcal{H}_i$ can be decomposed as,

$$\mathcal{H} = \bigotimes_{i=1}^n \mathcal{H}_i = \bigoplus_\lambda \big(\mathcal{M}_\lambda \otimes \mathcal{F}_\lambda\big), \qquad \text{where} \quad D_\lambda \equiv |\mathcal{M}_\lambda| = d_\lambda |G|^{n-1} d'^n = d_\lambda D/|G|, \tag{A.4}$$

where $D = |G|^n d'^n$ is the many-body Hilbert space dimension, and the symmetry operator $R_g = \bigotimes_{i=1}^n R_g^i$ acts as $R_g = \bigoplus_\lambda (\mathbb{1}_{\mathcal{M}_\lambda} \otimes R_g^\lambda)$. As a sanity check, one can easily compute $\sum_\lambda d_\lambda D_\lambda = (D/|G|)\sum_\lambda d_\lambda^2 = D$. If desired, one can also write the many-body Hilbert space as $\mathcal{H} \cong \mathbb{C}^{D/|G|} \otimes \mathcal{F}_{\text{reg}}$, where R_g acts in the regular representation on the second register. This shows that the symmetry operators remain orthogonal to one another on the many-body Hilbert space, $\text{tr}\big(R_g R_h^{-1}\big) = D\delta_{g,h}$, as one would expect.

Additional notation.—We conclude this introduction by setting up several definitions for the analysis that follows. We let P_λ denote the projector onto the irrep λ, i.e. $P_\lambda = \mathbb{1}_{\mathcal{M}_\lambda} \otimes \mathbb{1}_{\mathcal{F}_\lambda}$. We also define a set of computational basis states spanning each irrep Hilbert space, $\mathcal{F}_\lambda = \{|a\rangle \,|a = 1, \ldots, d_\lambda\}$. Any choice of basis is valid. We let $P_{k\ell}^\lambda$ denote the transition operator between basis states a and b, i.e. $P_{ab}^\lambda = \mathbb{1}_{\mathcal{M}_\lambda} \otimes |k\rangle\langle\ell|$. Both the set of $P_{k\ell}^\lambda/\sqrt{D_\lambda}$ and the set of $R_g/\sqrt{D}$ form complete orthonormal bases for the set of operators in $\text{span}\{R_g\}$ [111]. This implies the equality,

$$\sum_{\lambda k\ell} \frac{1}{D_\lambda} \text{tr}\Big((\cdot)P_{\ell k}^\lambda\Big) P_{k\ell}^\lambda = \frac{1}{D}\sum_g \text{tr}\big((\cdot)R_g^{-1}\big) R_g. \tag{A.5}$$

This will be useful for exchanging summations over the two sets of operators later on.

A.2 The exact symmetric Haar twirl

In this section, we derive an explicit and exact formula for the twirl over a symmetric Haar-random unitary. Our results apply to any unitary representation R of any symmetry group G on any Hilbert space $\mathcal{H}$.

Recall from the previous subsection [Eq. (A.2)] that any symmetric unitary U can be written in the form,

$$U = \bigoplus_\lambda \big(U_\lambda \otimes \mathbb{1}_{\mathcal{F}_\lambda}\big), \tag{A.6}$$

where $U_\lambda \in U(D_\lambda)$, the unitary group on D_λ states. The Haar measure on the group of symmetric unitaries is the unique measure that is invariant under left and right multiplication, $U \to UV$ and $U \to VU$, by any symmetric unitary V. One can easily see that this corresponds to drawing each U_λ independently from the Haar measure on $U(D_\lambda)$. We denote the resulting G-symmetric Haar ensemble as H_G.

We begin by providing an exact formula for the twirl over k copies of a symmetric Haar-random unitary. This exact formula is rather complicated; we will provide a much simpler approximate formula in the following subsection. To set up notation, we decompose the k-copy Hilbert space into a tensor product of k irreps, $\mathcal{H}^{\otimes k} = \bigoplus_{\boldsymbol{\lambda}} (\mathcal{M}_{\boldsymbol{\lambda}} \otimes \mathcal{F}_{\boldsymbol{\lambda}})$, where $\boldsymbol{\lambda} = (\lambda_1, \ldots, \lambda_k)$ is a length-k vector denoting the irrep in each of the k copies, and $\mathcal{M}_{\boldsymbol{\lambda}} = \bigotimes_{j=1}^k \mathcal{M}_{\lambda_j}$, $\mathcal{F}_{\boldsymbol{\lambda}} = \bigotimes_{j=1}^k \mathcal{F}_{\lambda_j}$. We define the basis operators $P_{\boldsymbol{k\ell}}^{\boldsymbol{\lambda}} = \bigotimes_{j=1}^k P_{k_j\ell_j}^{\lambda_j}$, where $\boldsymbol{k}, \boldsymbol{\ell}$ are also length-k vectors.

Proposition 1 (Exact symmetric Haar twirl)**.** *For any unitary representation of any symmetry group G. The twirl over k copies of a symmetric Haar-random unitary is,*

$$\Phi_H^G(\rho) \equiv \mathop{\mathbb{E}}_{U\sim H_G}\Big[U^{\otimes k}\rho U^{\dagger,\otimes k}\Big] = \sum_\sigma \sum_{\boldsymbol{\lambda k\ell}} \sum_{\tau_{\alpha_\lambda}} \text{tr}\Big(\rho\sigma^{-1}P_{\sigma(\boldsymbol{\ell})\boldsymbol{k}}^{\boldsymbol{\lambda}}\Big) \cdot \prod_\lambda \text{Wg}(\tau_{\alpha_\lambda}; D_\lambda) \cdot P_{\boldsymbol{k}\tau(\boldsymbol{\ell})}^{\boldsymbol{\lambda}}\,\tau, \tag{A.7}$$

for any $k \le \min_\lambda D_\lambda$. Here, the permutation $\tau \in S_k$ is defined as $\tau = \big(\bigotimes_{\alpha_\lambda} \tau_{\alpha_\lambda}\big)\sigma \in S_k$, where each $\tau_{\alpha_\lambda} \in S_{k_\lambda}$ permutes the copies $\alpha_\lambda = \{j|\lambda_j = \lambda\}$ that contain a given irrep λ, and $k_\lambda = |\alpha_\lambda|$.

We recall that $\text{Wg}(\tau_{\alpha_\lambda}; D_\lambda) = \text{Wg}_{\mathbb{1}_{\alpha_\lambda},\tau_{\alpha_\lambda}}(D_\lambda)$ denotes the Weingarten matrix element between τ_{α_λ} and the identity permutation on k_λ copies.

Proof. We write $\rho = \bigoplus_{\boldsymbol{\lambda},\boldsymbol{\lambda}'} \rho_{\boldsymbol{\lambda},\boldsymbol{\lambda}'}$ and $U = \bigoplus_{\boldsymbol{\lambda}} (U_{\boldsymbol{\lambda}} \otimes \mathbb{1}_{\mathcal{F}_{\boldsymbol{\lambda}}})$, where $U_{\boldsymbol{\lambda}} = \otimes_{i=1}^{k} U_{\lambda_i}$. We have,

$$\begin{aligned}
\Phi_H^G(\rho) &\equiv \bigoplus_{\boldsymbol{\lambda},\boldsymbol{\lambda}'} \mathop{\mathbb{E}}_{U_{\boldsymbol{\lambda}} \sim H} \Big[(U_{\boldsymbol{\lambda}} \otimes \mathbb{1}_{\mathcal{F}_{\boldsymbol{\lambda}}})\, \rho_{\boldsymbol{\lambda},\boldsymbol{\lambda}'}\, (U_{\boldsymbol{\lambda}'}^{\dagger} \otimes \mathbb{1}_{\mathcal{F}_{\boldsymbol{\lambda}'}}) \Big] \\
&\equiv \bigoplus_{\boldsymbol{\lambda},\boldsymbol{\lambda}'} \Big(\prod_{\lambda} \delta_{k_\lambda, k'_\lambda} \Big) \mathop{\mathbb{E}}_{U_{\boldsymbol{\lambda}} \sim H} \Big[(U_{\boldsymbol{\lambda}} \otimes \mathbb{1}_{\mathcal{F}_{\boldsymbol{\lambda}}})\, \rho_{\boldsymbol{\lambda},\boldsymbol{\lambda}'}\, (U_{\boldsymbol{\lambda}'}^{\dagger} \otimes \mathbb{1}_{\mathcal{F}_{\boldsymbol{\lambda}'}}) \Big] \\
&= \bigoplus_{\boldsymbol{\lambda},\boldsymbol{\lambda}'} \Big(\prod_{\lambda} \delta_{k_\lambda, k'_\lambda} \Big) \sum_{\sigma_{\alpha'_\lambda \to \alpha_\lambda}} \sum_{\tau_{\alpha'_\lambda \to \alpha_\lambda}} \mathrm{tr}_{\mathcal{M}_{\boldsymbol{\lambda}}}(\rho_{\boldsymbol{\lambda},\boldsymbol{\lambda}'}\, \sigma_{\mathcal{M}}^{-1}) \cdot \prod_{\lambda} \mathrm{Wg}_{\sigma_{\alpha'_\lambda \to \alpha_\lambda}, \tau_{\alpha'_\lambda \to \alpha_\lambda}}(D_\lambda) \cdot \tau_{\mathcal{M}} \\
&= \bigoplus_{\boldsymbol{\lambda}} \sum_{\sigma} \sum_{\tau_{\alpha_\lambda}} \mathrm{tr}_{\mathcal{M}_{\boldsymbol{\lambda}}}(\rho_{\boldsymbol{\lambda},\sigma^{-1}(\boldsymbol{\lambda})}\, \sigma_{\mathcal{M}}^{-1}) \cdot \prod_{\lambda} \mathrm{Wg}(\tau_{\alpha_\lambda}; D_\lambda) \cdot \tau_{\mathcal{M}} \\
&= \bigoplus_{\boldsymbol{\lambda}} \sum_{\sigma} \sum_{\tau_{\alpha_\lambda}} \mathrm{tr}\Big(\rho_{\boldsymbol{\lambda},\sigma^{-1}(\boldsymbol{\lambda})} \big[\sigma_{\mathcal{M}}^{-1} \otimes |\boldsymbol{\ell}\rangle\langle \boldsymbol{k}|_{\mathcal{F}_{\boldsymbol{\lambda}}} \big] \Big) \cdot \prod_{\lambda} \mathrm{Wg}(\tau_{\alpha_\lambda}; D_\lambda) \cdot \big[\tau_{\mathcal{M}} \otimes |\boldsymbol{k}\rangle\langle \boldsymbol{\ell}|_{\mathcal{F}_{\boldsymbol{\lambda}}} \big] \\
&= \bigoplus_{\boldsymbol{\lambda}} \sum_{\sigma} \sum_{\tau_{\alpha_\lambda}} \mathrm{tr}\Big(\rho_{\boldsymbol{\lambda},\sigma^{-1}(\boldsymbol{\lambda})}\, \sigma^{-1} \big[\mathbb{1}_{\mathcal{M}_{\boldsymbol{\lambda}}} \otimes |\sigma(\boldsymbol{\ell})\rangle\langle \boldsymbol{k}|_{\mathcal{F}_{\boldsymbol{\lambda}}} \big] \Big) \cdot \prod_{\lambda} \mathrm{Wg}(\tau_{\alpha_\lambda}; D_\lambda) \cdot \big[\mathbb{1}_{\mathcal{M}_{\boldsymbol{\lambda}}} \otimes |\boldsymbol{k}\rangle\langle \tau(\boldsymbol{\ell})|_{\mathcal{F}_{\boldsymbol{\lambda}}} \big] \tau \\
&= \sum_{\boldsymbol{\lambda} \boldsymbol{k} \boldsymbol{\ell}} \sum_{\sigma} \sum_{\tau_{\alpha_\lambda}} \mathrm{tr}\Big(\rho\, \sigma^{-1} P^{\boldsymbol{\lambda}}_{\sigma(\boldsymbol{\ell}),\boldsymbol{k}} \Big) \cdot \prod_{\lambda} \mathrm{Wg}(\tau_{\alpha_\lambda}; D_\lambda) \cdot P^{\boldsymbol{\lambda}}_{\boldsymbol{k},\tau(\boldsymbol{\ell})} \tau
\end{aligned} \tag{A.8}$$

In the second line, we note that the expectation value vanishes unless there are an equal number k_λ of U_λ as there are k'_λ of $U_\lambda^\dagger$ (where k'_λ denotes the number of copies with irrep λ in $\boldsymbol{\lambda}'$). In the third line, we compute the Haar average over k_λ copies for each irrep λ. The expression in terms of Weingarten functions is valid as long as $k_\lambda \leq D_\lambda$ for each λ; this is guaranteed when $k \leq \min_\lambda D_\lambda$. This produces a sum over permutations $\sigma = \bigotimes_\lambda \sigma_{\alpha'_\lambda \to \alpha_\lambda}$, $\tau = \bigotimes_\lambda \tau_{\alpha'_\lambda \to \alpha_\lambda}$, which map the copies α'_λ with irrep λ in $\boldsymbol{\lambda}'$ to the copies α_λ with irrep λ in $\boldsymbol{\lambda}$. The permutations act on the $\mathcal{M}$ registers; specifically, they are isometries between $\mathcal{M}_{\boldsymbol{\lambda}'}$ and $\mathcal{M}_{\boldsymbol{\lambda}}$. In the fourth line, we note that we can combine the sums over $\boldsymbol{\lambda}'$ (with $k'_\lambda = k_\lambda$) and $\sigma_{\alpha'_\lambda \to \alpha_\lambda}$ into a single summation over all $\sigma \in S_k$. We also re-label the summation over τ, using $\tau_{\alpha_\lambda} \equiv \tau_{\alpha'_\lambda \to \alpha_\lambda} (\sigma_{\alpha'_\lambda \to \alpha_\lambda})^{-1}$. In the fifth line, we insert the complete basis of operators on $\mathcal{F}_{\boldsymbol{\lambda}}$. In the sixth line, we insert the identity $\mathbb{1} = \sigma^{-1}\sigma$ on the $\mathcal{F}_{\boldsymbol{\lambda}}$ register, and similar for τ. In the final line, we re-write the expression using the operators $P^{\boldsymbol{\lambda}}_{\boldsymbol{k}\boldsymbol{\ell}}$. This concludes the proof. □

While our proof of Proposition 1 is direct and fairly short, it can be somewhat opaque due to the various tensor sums and products. To this end, we also present a second, more heuristic derivation of Proposition 1 below. Our method is inspired by a similar technique for deriving the Haar twirl without symmetries. In particular, we first enumerate the operators that commute with $U^{\otimes k}$ for every U. We then define $\Phi_H^G(\rho)$ as the map which projects ρ onto the span of these operators.

Loosely speaking, there are two classes of operators that commute with any symmetric $U^{\otimes k}$. First, the permutation operators $\sigma \in S_k$ commute with any tensor product unitary, symmetric or not. Second, since U is symmetric, any tensor product of local symmetry operators $R_{\boldsymbol{g}} = \otimes_{j=1}^{k} R_{g_j}$ also commutes with $U^{\otimes k}$. Any product of an operator from the first set and the second set also commutes with $U^{\otimes k}$. Hence, the commutant is equal to the set of all such products, $\{R_{\boldsymbol{g}}\, \sigma\}$. In what follows, it will be more convenient to replace the local symmetry operators R_g with the irrep basis elements $P^\lambda_{k\ell}$. This is valid because any R_g can be written as a linear combination of $P^\lambda_{k\ell}$ and vice versa. Thus, the commutant is equivalently spanned by the set $\{P^{\boldsymbol{\lambda}}_{\boldsymbol{k}\boldsymbol{\ell}}\, \sigma\}$.

To project an operator ρ onto the commutant, we need to know the overlaps of each operator in

the commutant. These are given by the following expression,

$$G_{\sigma,\tau}^{\boldsymbol{\lambda k \ell},\boldsymbol{\lambda' k' \ell'}} = \mathrm{tr}\Big(P_{\boldsymbol{k\ell}}^{\boldsymbol{\lambda}}\, \sigma\tau^{-1} P_{\boldsymbol{\ell' k'}}^{\boldsymbol{\lambda'}}\Big) = \delta_{\boldsymbol{\lambda},\boldsymbol{\lambda'}}\delta_{\boldsymbol{k},\boldsymbol{k'}} \cdot \delta_{\sigma^{-1}(\boldsymbol{\lambda}),\tau^{-1}(\boldsymbol{\lambda})}\delta_{\sigma^{-1}(\boldsymbol{\ell}),\tau^{-1}(\boldsymbol{\ell'})} \cdot \prod_{\lambda} G(\sigma_\lambda \tau_\lambda^{-1}; D_\lambda), \quad \text{(A.9)}$$

where $\sigma_\lambda, \tau_\lambda$ refer to the restriction of σ, τ to input and output indices in representation λ (i.e. σ_λ is a shorthand for $\sigma_{\sigma^{-1}(\alpha_\lambda)\to\alpha_\lambda}$ in the notation used in our proof of Proposition 1). When viewed as a matrix, the inverse of $G_{\sigma,\tau}^{\boldsymbol{\lambda k \ell},\boldsymbol{\lambda' k' \ell'}}$ is,

$$\mathrm{Wg}_{\sigma,\tau}^{\boldsymbol{\lambda k \ell},\boldsymbol{\lambda' k' \ell'}} = \delta_{\boldsymbol{\lambda},\boldsymbol{\lambda'}}\delta_{\boldsymbol{k},\boldsymbol{k'}} \cdot \delta_{\tau^{-1}(\boldsymbol{\lambda}),\sigma^{-1}(\boldsymbol{\lambda})}\delta_{\tau^{-1}(\boldsymbol{\ell}),\sigma^{-1}(\boldsymbol{\ell'})} \cdot \prod_{\lambda} \mathrm{Wg}(\sigma_\lambda \tau_\lambda^{-1}; D_\lambda). \quad \text{(A.10)}$$

We will verify this in the following paragraph. The inverse matrix allows us to immediately write an expression for the symmetric Haar twirl,

$$\Phi_H^G(\rho) = \sum_{\sigma,\tau} \sum_{\boldsymbol{\lambda k \ell}} \sum_{\boldsymbol{\lambda' k' \ell'}} \mathrm{tr}\Big(\rho\sigma^{-1} P_{\boldsymbol{\ell k}}^{\boldsymbol{\lambda}}\Big) \cdot \mathrm{Wg}_{\sigma,\tau}^{\boldsymbol{\lambda k \ell},\boldsymbol{\lambda' k' \ell'}} \cdot P_{\boldsymbol{k'\ell'}}^{\boldsymbol{\lambda'}}\, \tau. \quad \text{(A.11)}$$

One can see that $\Phi_H^G(\rho)$ is the unique map that projects ρ onto the span of the commutant. Inserting our expression for $\mathrm{Wg}_{\sigma,\tau}^{\boldsymbol{\lambda k \ell},\boldsymbol{\lambda' k' \ell'}}$ into the above, and re-indexing $\boldsymbol{\ell} \to \sigma(\boldsymbol{\ell})$, we find,

$$\Phi_H^G(\rho) = \sum_{\sigma,\tau} \sum_{\boldsymbol{\lambda k \ell}} \delta_{\tau^{-1}(\boldsymbol{\lambda}),\sigma^{-1}(\boldsymbol{\lambda})} \cdot \mathrm{tr}\Big(\rho\sigma^{-1} P_{\sigma(\boldsymbol{\ell})\boldsymbol{k}}^{\boldsymbol{\lambda}}\Big) \cdot \prod_{\lambda} \mathrm{Wg}(\sigma_\lambda \tau_\lambda^{-1}; D_\lambda) \cdot P_{\boldsymbol{k}\tau(\boldsymbol{\ell})}^{\boldsymbol{\lambda}}\, \tau. \quad \text{(A.12)}$$

We can then note that the delta function $\delta_{\tau^{-1}(\boldsymbol{\lambda}),\sigma^{-1}(\boldsymbol{\lambda})}$ amounts to restricting the sum over τ to those of the form $\tau = \big(\bigotimes_{\alpha_\lambda} \tau_{\alpha_\lambda}\big)\sigma \in S_k$ (as in Proposition 1). Implementing this substitution, i.e. replacing $\sum_\tau \delta_{\tau^{-1}(\boldsymbol{\lambda}),\sigma^{-1}(\boldsymbol{\lambda})} \to \sum_{\tau_{\alpha_\lambda}}$ and $\sigma_\lambda \tau_\lambda^{-1} \to \tau_{\alpha_\lambda}^{-1}$, the expression above reduces precisely to our desired formula, Eq. (A.7).

We conclude by showing that $\mathrm{Wg}_{\sigma,\tau}^{\boldsymbol{\lambda k \ell},\boldsymbol{\lambda' k' \ell'}}$ is indeed the inverse of $G_{\sigma,\tau}^{\boldsymbol{\lambda k \ell},\boldsymbol{\lambda' k' \ell'}}$. We compute,

$$\begin{aligned}
&\sum_{\tau} \sum_{\boldsymbol{\lambda' k' \ell'}} G_{\sigma,\tau}^{\boldsymbol{\lambda k \ell},\boldsymbol{\lambda' k' \ell'}}\, \mathrm{Wg}_{\tau,\pi}^{\boldsymbol{\lambda' k' \ell'},\boldsymbol{\lambda'' k'' \ell''}} \\
&\qquad = \delta_{\boldsymbol{\lambda},\boldsymbol{\lambda''}}\delta_{\boldsymbol{k},\boldsymbol{k''}}\delta_{\sigma^{-1}(\boldsymbol{\ell}),\pi^{-1}(\boldsymbol{\ell''})}\delta_{\sigma^{-1}(\boldsymbol{\lambda}),\pi^{-1}(\boldsymbol{\lambda})} \cdot \sum_{\tau} \delta_{\sigma^{-1}(\boldsymbol{\lambda}),\tau^{-1}(\boldsymbol{\lambda})} \cdot \prod_{\lambda} G_{\sigma_\lambda,\tau_\lambda}(D_\lambda)\, \mathrm{Wg}_{\tau_\lambda,\pi_\lambda}(D_\lambda) \\
&\qquad = \delta_{\boldsymbol{\lambda},\boldsymbol{\lambda''}}\delta_{\boldsymbol{k},\boldsymbol{k''}}\delta_{\sigma^{-1}(\boldsymbol{\ell}),\pi^{-1}(\boldsymbol{\ell''})}\delta_{\sigma^{-1}(\boldsymbol{\lambda}),\pi^{-1}(\boldsymbol{\lambda})} \cdot \prod_{\lambda} \Big(\sum_{\tau_\lambda} G_{\sigma_\lambda,\tau_\lambda}(D_\lambda)\, \mathrm{Wg}_{\tau_\lambda,\pi_\lambda}(D_\lambda) \Big) \\
&\qquad = \delta_{\boldsymbol{\lambda},\boldsymbol{\lambda''}}\delta_{\boldsymbol{k},\boldsymbol{k''}}\delta_{\sigma^{-1}(\boldsymbol{\ell}),\pi^{-1}(\boldsymbol{\ell''})}\delta_{\sigma^{-1}(\boldsymbol{\lambda}),\pi^{-1}(\boldsymbol{\lambda})} \cdot \prod_{\lambda} \delta_{\sigma_\lambda,\pi_\lambda} \\
&\qquad = \delta_{\boldsymbol{\lambda},\boldsymbol{\lambda''}}\delta_{\boldsymbol{k},\boldsymbol{k''}}\delta_{\boldsymbol{\ell},\boldsymbol{\ell''}}\delta_{\sigma,\pi}.
\end{aligned} \quad \text{(A.13)}$$

In the first line, we note the delta functions in Eq. (A.9) and Eq. (A.10) enforce $\boldsymbol{\lambda} = \boldsymbol{\lambda'} = \boldsymbol{\lambda''}$, $\boldsymbol{k} = \boldsymbol{k'} = \boldsymbol{k''}$, and $\sigma^{-1}(\boldsymbol{\ell}) = \tau^{-1}(\boldsymbol{\ell'}) = \pi^{-1}(\boldsymbol{\ell''})$. This allows us to compute the sums over $\boldsymbol{\lambda'}, \boldsymbol{k'}, \boldsymbol{\ell'}$, which eliminates the delta functions involving these variables. In the second line, the condition $\tau^{-1}(\boldsymbol{\lambda}) = \sigma^{-1}(\boldsymbol{\lambda})$ restricts the sum over τ to permutations such that τ^{-1} maps each copy i with an irrep λ_i to a copy j such that $\lambda_{\sigma(j)} = \lambda_i$. In other words, the permutation τ maps each subset of input copies with an irrep λ, $\{i : \lambda_i = \lambda\}$, to a fixed subset of output copies $\{j : \lambda_{\sigma(j)} = \lambda\}$. Thus, the sum over permutations τ reduces to a product of individual sums, over permutations $\tau_\lambda \in S_{k_\lambda}$ that map $\{i : \lambda_i = \lambda\}$ to $\{j : \lambda_{\sigma(j)} = \lambda\}$, where $k_\lambda = |\{i : \lambda_i = \lambda\}| = |\{j : \lambda_{\sigma(j)} = \lambda\}|$. In the third and fourth line, we find that each sum over τ_λ produces a delta function $\delta_{\sigma_\lambda,\pi_\lambda}$, which causes the remaining expression to reduce to the identity matrix. This completes our alternate derivation of Proposition 1.

A.3 Approximate symmetric Haar twirl

The expression for the exact symmetric Haar twirl is somewhat daunting, owing to the numerous summations and Weingarten matrix elements. In this section, we show that a much simpler formula suffices when the Hilbert space dimension is large. Our approximate formula holds for any discrete symmetry group G acting in the regular representation on the Hilbert space.

Lemma 2 (Approximate symmetric Haar twirl)**.** *For any discrete on-site symmetry group G. The symmetric Haar twirl Φ_H^G is approximated by the map,*

$$\Phi_a^G(\rho) = \frac{1}{D^k} \sum_{\sigma} \sum_{\boldsymbol{g}} \mathrm{tr}\big(\rho\sigma^{-1} R_{\boldsymbol{g}}^{-1}\big)\, R_{\boldsymbol{g}}\, \sigma, \tag{A.14}$$

up to relative error, $(1-\varepsilon)\Phi_H^G \preceq \Phi_{\mathcal{E}}^G \preceq (1+\varepsilon)\Phi_H^G$, with $\varepsilon = |G|k^2/D$ for any $k^2 \leq D/|G|$. Here, $\mathcal{A} \preceq \mathcal{B}$ denotes that $\mathcal{B}-\mathcal{A}$ is a completely positive map.

We remark that the approximate symmetric Haar twirl can be equivalently written in terms of the irreducible subspace projectors, $P_{\boldsymbol{k\ell}}^{\boldsymbol{\lambda}}$, as,

$$\Phi_a^G(\rho) = \sum_{\sigma} \sum_{\boldsymbol{\lambda k\ell}} \mathrm{tr}\Big(\rho\sigma^{-1} P_{\boldsymbol{\ell k}}^{\boldsymbol{\lambda}}\Big) \cdot \frac{1}{\prod_\lambda D_\lambda^{k_\lambda}} \cdot P_{\boldsymbol{k\ell}}^{\boldsymbol{\lambda}}\, \sigma. \tag{A.15}$$

This follows from writing $R_{\boldsymbol{g}} = \bigotimes_{i=1}^n R_{g_i}$ and $P_{\boldsymbol{k\ell}}^{\boldsymbol{\lambda}} = \bigotimes_{i=1}^n P_{k_i\ell_i}^{\lambda_i}$, and noting that within each copy, we have $\frac{1}{D}\sum_g \mathrm{tr}\big((\cdot)R_g^{-1}\big)R_g = \sum_{\lambda k\ell} \frac{1}{D_\lambda}\mathrm{tr}\big((\cdot)P_{\ell k}^\lambda\big)P_{k\ell}^\lambda$ from Eq. (A.5).

We will break our proof of Lemma 2 into two parts. First, we provide a general lemma for bounding the relative error of any random unitary ensemble with respect to the approximate symmetric Haar twirl, Φ_a^G. This lemma will also be used in our proof of the gluing lemma for symmetric unitaries (Lemma 4) later on. Second, we apply this lemma to prove Lemma 2.

A.3.1 Bounding the relative error via the additive error on the EPR state

In practice, it is usually much easier to bound the additive error of a random unitary ensemble rather than the relative error. To this end, we provide a helpful lemma for converting from additive to relative error, with respect to the approximate symmetric Haar twirl, Φ_a^G.

Lemma 3 (Relative error from EPR states)**.** *For any discrete on-site symmetry group G. Let $\mathcal{E}$ be any symmetric unitary ensemble, and $\Phi_{\mathcal{E}}$ its twirl on $\mathcal{H}^{\otimes k}$. The twirl is approximated by Φ_a^G up to relative error, $(1-\varepsilon)\Phi_a^G \preceq \Phi_{\mathcal{E}} \preceq (1+\varepsilon)\Phi_a^G$, where*

$$\varepsilon = \frac{k!|G|^k}{D^{2k}} \big\| \big[\delta\Phi \otimes \mathbb{1}\big](P_{\mathrm{EPR}}) \big\|_\infty. \tag{A.16}$$

Here, $\delta\Phi \equiv \Phi_{\mathcal{E}} - \Phi_a^G$, and P_{EPR} is the projector onto the EPR state on $\mathcal{H}^{\otimes k} \otimes \mathcal{H}^{\otimes k}$.

Our proof closely resembles that of Lemma 2 in Ref. [30] for random unitaries without symmetries.

Proof. Let $\rho_a^G \equiv [\Phi_a^G \otimes \mathbb{1}](P_{\mathrm{EPR}})$ and $\rho_{\mathcal{E}} \equiv [\Phi_{\mathcal{E}} \otimes \mathbb{1}](P_{\mathrm{EPR}})$. Our proof proceeds in four steps.

(1) ρ_a^G is equal to $k!|G|^k/D^{2k}$ times a projector,

$$P \equiv \frac{1}{k!|G|^k} \sum_{\sigma} \sum_{\boldsymbol{g}} R_{\boldsymbol{g}}\, \sigma \otimes R_{\boldsymbol{g}}^*\, \sigma \tag{A.17}$$

Hence, on the projected subspace, ρ_a^G has a flat spectrum with eigenvalue $k!|G|^k/D^{2k}$. To show that P is a projector, we compute that it squares to one,

$$\begin{aligned} P^2 &= \frac{1}{k!^2|G|^{2k}} \sum_{\sigma,\sigma'} \sum_{\boldsymbol{g},\boldsymbol{g}'} R_{\boldsymbol{g}}\, \sigma\, R_{\boldsymbol{g}'}\, \sigma' \otimes R^*_{\boldsymbol{g}}\, \sigma\, R^*_{\boldsymbol{g}'}\, \sigma' \\ &= \frac{1}{k!^2|G|^{2k}} \sum_{\sigma,\sigma'} \sum_{\boldsymbol{g},\boldsymbol{g}'} R_{\boldsymbol{g}\cdot\sigma(\boldsymbol{g}')}\, \sigma\sigma' \otimes R^*_{\boldsymbol{g}\cdot\sigma(\boldsymbol{g}')}\, \sigma\sigma' \\ &= \frac{1}{k!^2|G|^{2k}} \sum_{\sigma,\sigma'} \sum_{\boldsymbol{g},\boldsymbol{g}'} R_{\boldsymbol{g}\cdot\boldsymbol{g}'}\, \sigma\sigma' \otimes R^*_{\boldsymbol{g}\cdot\boldsymbol{g}'}\, \sigma\sigma' \\ &= \frac{1}{k!|G|^k} \sum_{\sigma} \sum_{\boldsymbol{g}} R_{\boldsymbol{g}}\, \sigma \otimes R^*_{\boldsymbol{g}}\, \sigma = P. \end{aligned} \tag{A.18}$$

In the second line, we re-index $\boldsymbol{g}' \to \sigma^{-1}(\boldsymbol{g}')$. In the fourth line, we compute the sums over σ' and $\boldsymbol{g}'$ by re-indexing $\sigma \to \sigma\sigma'^{-1}$ and $\boldsymbol{g} \to \boldsymbol{g}\boldsymbol{g}'^{-1}$. The sums then give $\sum_{\sigma'} 1 = k!$ and $\sum_{\boldsymbol{g}'} 1 = |G|^k$.

(2) $\rho_{\mathcal{E}}$ has support entirely within the support of P. This follows because $P(U^{\otimes k} \otimes \mathbb{1})\,|\Psi_{\mathrm{EPR}}\rangle = (U^{\otimes k} \otimes \mathbb{1})P\,|\Psi_{\mathrm{EPR}}\rangle = (U^{\otimes k} \otimes \mathbb{1})\,|\Psi_{\mathrm{EPR}}\rangle$ for any $U \sim \mathcal{E}$, where $|\Psi_{\mathrm{EPR}}\rangle\langle\Psi_{\mathrm{EPR}}| \equiv P_{\mathrm{EPR}}$. The first equality follows because any symmetric tensor product unitary $U^{\otimes k}$ commutes with any permutation σ and any tensor product $R_{\boldsymbol{g}}$ of symmetry operators. The second equality follows because $(R_{\boldsymbol{g}}\sigma \otimes R^*_{\boldsymbol{g}}\sigma)\,|\Psi_{\mathrm{EPR}}\rangle = (R_{\boldsymbol{g}}\sigma\sigma^{-1}R_{\boldsymbol{g}}^{-1} \otimes \mathbb{1})\,|\Psi_{\mathrm{EPR}}\rangle = |\Psi_{\mathrm{EPR}}\rangle$.

(3) Steps (1) and (2) immediately imply that the twirl has relative error ε [Eq. (A.16)] on the EPR state.

(4) The relative error on the EPR state upper bounds the relative error on any state. This follows because we can express $\Phi(\rho) = D^{2k}\,\mathrm{tr}_2((\mathbb{1} \otimes \rho^T)[\Phi \otimes \mathbb{1}](P_{\mathrm{EPR}}))$ for any Φ, where the trace is over the second copy of $\mathcal{H}^{\otimes k}$. □

A.3.2 Proof of Lemma 2

With Lemma 3 in hand, we will now prove Lemma 2. Let $\rho_a^G = [\Phi_a^G \otimes \mathbb{1}](P_{\mathrm{EPR}})$ and $\rho_H^G = [\Phi_H^G \otimes \mathbb{1}](P_{\mathrm{EPR}})$ denote the approximate and exact Haar twirls applied to the EPR state. For the purposes of this proof, it will be convenient to write ρ_H^G by inserting the EPR state $\rho \to P_{\mathrm{EPR}}$ into the third line of Eq. (A.8). To do so, we first decompose the EPR state into a tensor sum of EPR states on each irreducible representation,

$$P_{\mathrm{EPR}} = \bigoplus_{\boldsymbol{\lambda},\boldsymbol{\lambda}'} \left(|\Psi^{\mathcal{M}_{\boldsymbol{\lambda}}}_{\mathrm{EPR}}\rangle\langle\Psi^{\mathcal{M}_{\boldsymbol{\lambda}'}}_{\mathrm{EPR}}| \otimes |\Psi^{\mathcal{F}_{\boldsymbol{\lambda}}}_{\mathrm{EPR}}\rangle\langle\Psi^{\mathcal{F}_{\boldsymbol{\lambda}'}}_{\mathrm{EPR}}| \right) \cdot \prod_{\lambda} (d_\lambda D_\lambda/D)^{k_\lambda/2} \cdot \prod_{\lambda'} (d_{\lambda'} D_{\lambda'}/D)^{k'_{\lambda'}/2}. \tag{A.19}$$

Here, $\boldsymbol{\lambda}$ labels the irrep of the ket and $\boldsymbol{\lambda}'$ of the bra, and $k_\lambda, k'_{\lambda'}$ indicate the multiplicity of an irrep λ, λ' in $\boldsymbol{\lambda}, \boldsymbol{\lambda}'$. The multiplicative factors at the right ensure that each irrep has trace equal to its dimension $d_\lambda D_\lambda$ divided by the total Hilbert space dimension D (since the EPR state has flat

support over the total Hilbert space). This yields an expression for ρ_H^G,

$$\begin{aligned}
\rho_H^G &= \bigoplus_{\boldsymbol{\lambda}} \sum_{\sigma,\tau_{\alpha_\lambda}} \mathrm{tr}_{\mathcal{M}_{\boldsymbol{\lambda}}}([P_{\mathrm{EPR}}]_{\boldsymbol{\lambda},\sigma^{-1}(\boldsymbol{\lambda})}\,[\sigma^{-1}]_{\mathcal{M}_{\sigma^{-1}(\boldsymbol{\lambda})}\to\mathcal{M}_{\boldsymbol{\lambda}}}) \cdot \prod_\lambda \mathrm{Wg}(\tau_{\alpha_\lambda}; D_\lambda) \cdot [\tau]_{\mathcal{M}_{\boldsymbol{\lambda}}\to\mathcal{M}_{\sigma^{-1}(\boldsymbol{\lambda})}}, \\
&= \frac{|G|^k}{D^{2k}} \bigoplus_{\boldsymbol{\lambda}} \sum_{\sigma,\tau_{\alpha_\lambda}} \left(\prod_\lambda \left(D_\lambda^{k_\lambda} \,\mathrm{Wg}(\tau_{\alpha_\lambda}; D_\lambda) \right) \cdot \left([\tau\otimes\sigma]_{\mathcal{M}_{\boldsymbol{\lambda}}^{\otimes 2}\to\mathcal{M}_{\sigma^{-1}(\boldsymbol{\lambda})}^{\otimes 2}} \otimes |\Psi_{\mathrm{EPR}}^{\mathcal{F}_{\boldsymbol{\lambda}}}\rangle\langle\Psi_{\mathrm{EPR}}^{\mathcal{F}_{\sigma^{-1}(\boldsymbol{\lambda})}}| \right) \right) \\
&= \frac{|G|^k}{D^{2k}} \sum_\sigma \left(\bigoplus_{\boldsymbol{\lambda}} \left(\sum_{\tau_{\alpha_\lambda}} \prod_\lambda \left(D_\lambda^{k_\lambda} \,\mathrm{Wg}(\tau_{\alpha_\lambda}; D_\lambda) \right) \cdot \left([(\tau\sigma^{-1})\otimes \mathbb{1}]_{\mathcal{M}_{\boldsymbol{\lambda}}^{\otimes 2}} \otimes |\Psi_{\mathrm{EPR}}^{\mathcal{F}_{\boldsymbol{\lambda}}}\rangle\langle\Psi_{\mathrm{EPR}}^{\mathcal{F}_{\boldsymbol{\lambda}}}| \right) \right) \right) (\sigma\otimes\sigma).
\end{aligned} \tag{A.20}$$

We add subscripts to the various operators to make it clear which subspaces they map between. In the second line, we apply the identity, $\mathrm{tr}_{\mathcal{M}_{\boldsymbol{\lambda}}}([P_{\mathrm{EPR}}]_{\boldsymbol{\lambda},\sigma^{-1}(\boldsymbol{\lambda})}\,[\sigma^{-1}]_{\mathcal{M}_{\sigma^{-1}(\boldsymbol{\lambda})}\to\mathcal{M}_{\boldsymbol{\lambda}}}) = \prod_\lambda (1/D_\lambda)^{k_\lambda} \cdot [\sigma]_{\mathcal{M}_{\boldsymbol{\lambda}}\to\mathcal{M}_{\sigma^{-1}(\boldsymbol{\lambda})}}$, to take the partial trace over $\mathcal{M}_{\boldsymbol{\lambda}}$ on the left side of the the EPR projector. This follows from the standard identity for the partial trace over the left half of an EPR pair, $\mathrm{tr}_1([A\otimes\mathbb{1}]P_{\mathrm{EPR}}) = A^T/D$. We also use $\prod_\lambda (d_\lambda D_\lambda/D)^{k_\lambda} \cdot \prod_\lambda (1/D_\lambda)^{k_\lambda} = (|G|^k/D^k)\prod_\lambda D_\lambda^{k_\lambda}$ to simplify the constant factors arising from Eq. (A.19) and this expression. In the final line, we pull out a factor $\sigma\otimes\sigma$, which will be convenient later on.

We can obtain an analogous expression for the EPR state under the approximate Haar twirl, by replacing $\prod_\lambda D_\lambda^{k_\lambda}\,\mathrm{Wg}(\tau_{\alpha_\lambda}; D_\lambda) \to \delta_{\sigma,\tau}$. (To be precise, this follows from Eq. (A.15) and manipulations exactly analogous to the final four lines of Eq. (A.8).) This yields,

$$\rho_a^G = \frac{|G|^k}{D^{2k}} \sum_\sigma \left(\bigoplus_{\boldsymbol{\lambda}} \left(\mathbb{1}_{\mathcal{M}_{\boldsymbol{\lambda}}^{\otimes 2}} \otimes |\Psi_{\mathrm{EPR}}^{\mathcal{F}_{\boldsymbol{\lambda}}}\rangle\langle\Psi_{\mathrm{EPR}}^{\mathcal{F}_{\boldsymbol{\lambda}}}| \right) \right) (\sigma\otimes\sigma). \tag{A.21}$$

Taking the difference between the two states and applying the triangle inequality, with $\|\bigoplus_{\boldsymbol{\lambda}} A_{\boldsymbol{\lambda}}\|_\infty = \max_{\boldsymbol{\lambda}} \|A_{\boldsymbol{\lambda}}\|_\infty$, we find

$$\begin{aligned}
\|\rho_H^G - \rho_a^G\|_\infty &\leq \frac{|G|^k}{D^{2k}} \sum_\sigma \max_{\boldsymbol{\lambda}} \left(\sum_{\tau_{\alpha_\lambda}} \left| \prod_\lambda D_\lambda^{k_\lambda} \,\mathrm{Wg}(\tau_{\alpha_\lambda}; D_\lambda) - \delta_{\sigma,\tau} \right| \right) \\
&= \frac{|G|^k}{D^{2k}} \sum_\sigma \max_{\boldsymbol{\lambda}} \left(\sum_{\tau_{\alpha_\lambda}} \left(\prod_\lambda D_\lambda^{k_\lambda} |\mathrm{Wg}(\tau_{\alpha_\lambda}; D_\lambda)| - \delta_{\sigma,\tau} \right) \right) \\
&= \frac{|G|^k}{D^{2k}} \sum_\sigma \max_{\boldsymbol{\lambda}} \left(\prod_\lambda D_\lambda^{k_\lambda} \frac{(D_\lambda - k_\lambda)!}{D_\lambda!} - 1 \right) \\
&\leq \frac{|G|^k}{D^{2k}} \sum_\sigma \max_{\boldsymbol{\lambda}} \left(\prod_\lambda (1 + k_\lambda^2/D_\lambda) - 1 \right) \\
&= \frac{|G|^k}{D^{2k}} \sum_\sigma \frac{k^2}{\min_\lambda D_\lambda} = \frac{k!|G|^k}{D^{2k}} \frac{|G|k^2}{D}.
\end{aligned} \tag{A.22}$$

In the second line, we use that the diagonal Weingarten elements are greater than one, $D_\lambda^{k_\lambda}\,\mathrm{Wg}(\mathbb{1}_{\alpha_\lambda}; D_\lambda) \geq 1$, and in the third line, we use that $\sum_{\tau_{\alpha_\lambda}} |\mathrm{Wg}(\tau_{\alpha_\lambda}; D_\lambda)| = (D_\lambda - k_\lambda)!/D_\lambda! \leq 1 + k_\lambda^2/D_\lambda$ [57]. In the final line, we note that the maximum is achieved when $\boldsymbol{\lambda}$ contains k copies of the same value of λ, for λ such that D_λ is minimal. We also have $\min_\lambda D_\lambda = \min_\lambda d_\lambda D/|G| \geq D/|G|$. Applying Lemma 3 completes our proof. □

A.4 Symmetric unitary designs

In many cases, one would like quantify whether a given symmetric random unitary ensemble is "close" to the symmetric Haar ensemble. One way to do so is through the notion of a symmetric unitary k-design. We say that an ensemble of symmetric unitaries forms an ε-approximate symmetric unitary k-design if it reproduces the k-th moment of the symmetric Haar ensemble up to relative error ε.

Definition 1 (Approximate symmetric unitary designs)**.** *For any symmetry group G. A symmetric unitary ensemble $\mathcal{E}$ is an ε-approximate symmetric unitary k-design if*

$$(1-\varepsilon)\,\Phi_H^G \preceq \Phi_{\mathcal{E}} \preceq (1+\varepsilon)\,\Phi_H^G, \tag{A.23}$$

where the quantum channel $\Phi_{\mathcal{E}}(\cdot)$ is defined via

$$\Phi_{\mathcal{E}}(A) \coloneqq \mathop{\mathbb{E}}_{U\sim\mathcal{E}}\left[U^{\otimes k} A U^{\dagger,\otimes k}\right], \tag{A.24}$$

and similarly for the symmetric Haar ensemble.

We will use the notion of symmetric unitary designs to state our gluing lemma in the following subsection.

A.5 Proof of Lemma 1: Gluing symmetric random unitaries

We now arrive at one of our main results on symmetric random unitaries, the symmetric gluing lemma (Lemma 1). The lemma shows that small symmetric random unitaries can "glue" together to form larger symmetric random unitaries in much the same manner as random unitaries without symmetries [30]. Let first state a formal version of Lemma 1 with tighter error bounds.

Lemma 4 (Gluing two symmetric random unitaries; formal version of Lemma 1)**.** *For any discrete on-site symmetry group G. Let A, B, C be three disjoint subsystems. Consider a symmetric random unitary given by $V_{ABC} = U_{AB}U_{BC}$, where U_{AB} and U_{BC} are drawn from ε_{AB} and ε_{BC}-approximate symmetric unitary k-designs, respectively. Then V_{ABC} is an ε-approximate symmetric unitary k-design with*

$$1+\varepsilon = (1+\varepsilon_{AB})(1+\varepsilon_{BC})\left(1-\frac{k^2|G|}{2D_{AB}}\right)^{-1}\left(1-\frac{k^2|G|}{2D_{BC}}\right)^{-1} e^{k(k-1)|G|/2D_B}\left(1+\frac{k^2|G|}{D_{ABC}}\right), \tag{A.25}$$

for any $k^2 \le D_B/|G|$. Here, $D_\alpha = 2^{|\alpha|}$ is the Hilbert space dimension of subsystem α.

Our proof of the gluing lemma utilizes the simple approximate expression for the symmetric Haar twirl derived in the previous subsections.

From the gluing lemma, we can immediately show that symmetric two-layer symmetric circuit forms a unitary k-design.

Theorem 7 (Gluing small symmetric random unitaries)**.** *For any discrete on-site symmetry group and any approximation error $\varepsilon \le 1$. Suppose each small random unitary in the two-layer brickwork ensemble $\mathcal{E}$ is drawn from an $\frac{\varepsilon}{n}$-approximate symmetric unitary k-design on 2ξ qubits with circuit depth d. Then $\mathcal{E}$ forms an ε-approximate symmetric unitary k-design on n qubits with depth $2d$, whenever the local patch size satisfies $\xi \ge \log_{|G|}(nk^2|G|/\varepsilon)$.*

This will form the basis of our proof of Theorem 3, on low-depth symmetric PRUs, in the next section.

In what follows, we first prove the symmetric gluing lemma and then the two-layer circuit theorem.

Proof of Lemma 4. Let $\delta_{AB} = k^2|G|/2D_{AB}/(1-k^2|G|/2D_{AB})$, $\delta_{BC} = k^2|G|/2D_{BC}/(1-k^2|G|/2D_{BC})$; elementary algebra gives $1+\delta_{AB} = (1-k^2|G|/2D_{AB})^{-1}$ and similar for δ_{BC}. From Lemma 2, we can approximate the twirl over V_{ABC} by $(\Phi_a^G)_{AB} \circ (\Phi_a^G)_{BC}$ up to relative error $(1+\varepsilon_{AB})(1+\delta_{AB})(1+\varepsilon_{BC})(1+\delta_{BC}) - 1$. To proceed, let

$$\rho_a^G = [\Phi_a^G \otimes \mathbb{1}](P_{\mathrm{EPR}}) = \frac{1}{D^{2k}} \sum_{\sigma} \sum_{\boldsymbol{g}} R_{\boldsymbol{g}}\, \sigma \otimes R_{\boldsymbol{g}}^*\, \sigma, \tag{A.26}$$

and

$$\begin{aligned} \rho_{\mathcal{E},a} &\equiv \left[\left((\Phi_a^G)_{AB} \circ (\Phi_a^G)_{BC}\right) \otimes \mathbb{1}\right](P_{\mathrm{EPR}}) \\ &= \frac{1}{D^k} \frac{1}{D_{AB}^k D_{BC}^k} \sum_{\sigma\tau} \sum_{\boldsymbol{g}\boldsymbol{g}'} G_{\sigma,\tau}^{\boldsymbol{g},\boldsymbol{g}'} \cdot [(R_{\boldsymbol{g}}\sigma)_{AB} \otimes (R_{\boldsymbol{g}'}\tau)_C] \otimes [(R_{\boldsymbol{g}}^*\sigma)_A \otimes (R_{\boldsymbol{g}'}^*\tau)_{BC}], \end{aligned} \tag{A.27}$$

where we denote $G_{\sigma,\tau}^{\boldsymbol{g},\boldsymbol{g}'} = \mathrm{tr}_B((R_{\boldsymbol{g}}\sigma\tau^{-1}R_{\boldsymbol{g}'}^{-1})_B)$. We note that the diagonal terms in $\rho_{\mathcal{E},a}$, with $\sigma = \tau$, have $G_{\sigma,\sigma}^{\boldsymbol{g},\boldsymbol{g}'} = \prod_{i=1}^k \mathrm{tr}_{B_i}(R_{g_i}R_{g_i'}^{-1}) = D_B^k\, \delta_{\boldsymbol{g},\boldsymbol{g}'}$, and thus yield precisely the expression for ρ_a^G. Applying the triangle inequality to the remaining terms, which have $\sigma \neq \tau$, we find

$$\begin{aligned} \|\rho_a^G - \rho_{\mathcal{E},a}\|_\infty &\leq \frac{1}{D^k} \frac{1}{D_{AB}^k D_{BC}^k} \sum_{\sigma\neq\tau} \sum_{\boldsymbol{g}\boldsymbol{g}'} \left|G_{\sigma,\tau}^{\boldsymbol{g},\boldsymbol{g}'}\right| \\ &= \frac{1}{D^k} \frac{1}{D_{AB}^k D_{BC}^k} \sum_{\sigma\neq\tau} \sum_{\boldsymbol{g}\boldsymbol{g}'} \left| \prod_{\ell\in\sigma\tau^{-1}} \mathrm{tr}_B\left(\prod_{i\in\ell} R_{g_i'}^{-1} R_{g_i}\right)\right| \\ &\leq \frac{1}{D^k} \frac{1}{D_{AB}^k D_{BC}^k} \sum_{\sigma\neq\tau} \sum_{\boldsymbol{g}\boldsymbol{g}'} D_B^{k-|\sigma\tau^{-1}|} \prod_{\ell\in\sigma\tau^{-1}} \delta\Big(\prod_{i\in\ell} g_i' g_i^{-1} = e\Big) \\ &= \frac{1}{D^{2k}} \sum_{\sigma\neq\tau} D_B^{-|\sigma\tau^{-1}|} |G|^{k+|\sigma\tau^{-1}|} \\ &= \frac{k!|G|^k}{D^{2k}} \sum_{\sigma\neq\mathbb{1}} (|G|/D_B)^{|\sigma|} \leq \frac{k!|G|^k}{D^{2k}} \left(e^{k(k-1)|G|/2D_B} - 1\right). \end{aligned} \tag{A.28}$$

In the second line, we write $G_{\sigma,\tau}^{\boldsymbol{g},\boldsymbol{g}'}$ as a product of traces over each cycle ℓ in $\sigma\tau^{-1}$. In the third line, we note that the trace over each cycle is zero unless the product of group elements appearing the cycle is the identity. Moreover, when this condition is satisfied, the trace is equal to D_B. The total number of cycles is $\#(\sigma\tau^{-1}) = k - |\sigma\tau^{-1}|$. In the fourth line, we compute the sum over $\boldsymbol{g}, \boldsymbol{g}'$. This yields $|G|^{2k-\#(\sigma\tau^{-1})} = |G|^{k+|\sigma\tau^{-1}|}$, where $\#(\sigma\tau^{-1})$ counts the number of delta functions that restrict the sum. In the final line, we use the well-known bound $\sum_\sigma D^{-|\sigma|} \leq e^{k(k-1)/2D}$ for $D \to D_B/|G|$ [30, 112].

We have $\|\rho_{\mathcal{E}} - \rho_a^G\|_\infty \leq (1+\varepsilon_{AB})(1+\delta_{AB})(1+\varepsilon_{BC})(1+\delta_{BC})\|\rho_{\mathcal{E},a} - \rho_a^G\|_\infty$, since $\rho_{\mathcal{E},a}$ approximates $\rho_{\mathcal{E}}$ up to relative error. Applying Lemma 3 (which bounds the relative error between $\Phi_{\mathcal{E}}$ and Φ_a^G) and Lemma 2 (which bounds the relative error between Φ_a^G and Φ_H^G) completes the proof. □

Proof of Theorem 7. Now that the symmetric gluing lemma is established, our proof exactly mirrors that of the proof of Theorem 1 in Ref. [30]. We apply Lemma 4 patch-by-patch $m-2$ times, where $m = n/\xi$ is the number of local patches. After all $m-2$ applications, one finds that the two-layer ensemble forms an approximate symmetric unitary k-design with error

$$(1+\varepsilon/n)^{m-1} \cdot \left(\left(1-k^2|G|/2q^2\right)^{-1} \left(1-k^2|G|/2q^2\right)^{-1} e^{k(k-1)|G|/2q} \left(1+k^2|G|/q^3\right)\right)^{m-2} - 1, \tag{A.29}$$

where we let $q \equiv |G|^\xi$ denote the Hilbert space dimension of a single patch. We also use $D_B \geq q$, $D_{AB}, D_{BC} \geq q^2$, and $D_{ABC} \geq q^3$. Using the inequalities $1/(1-x/2) \leq 1+x$ and $1+x \leq e^x$ several times, the above expression is upper bounded by,

$$(\ldots) \leq e^{(m-1)\varepsilon/n+(m-2)f(k,q)} - 1 \leq \frac{1}{\log(2)}\left((m-1)\varepsilon/n + (m-2)f(k,q)\right). \tag{A.30}$$

The second inequality follows from $e^x - 1 \leq x/\log(2)$ for $x \leq \log(2)$, and we abbreviate,

$$f(k,q) = \frac{k(k-1)|G|}{2q} + 2\frac{k^2|G|}{q^2} + \frac{k^2|G|}{q^3}, \tag{A.31}$$

which arise from the four error terms in the gluing lemma.

We would like to show that the total error, Eq. (A.29), is less than ε. We take $k \geq 2$ and $n \geq 3\xi$ since otherwise the theorem holds trivially. By assumption, we have $q \geq nk^2|G|/\varepsilon$. Combined with $k \geq 2$, $n \geq 3\xi$, $|G| \geq 1$, and $\varepsilon \leq 1$, this implies that $\xi \geq 5$ and so $q \geq 32$. The first term in Eq. (A.29) is therefore less than,

$$\frac{(m-1)\varepsilon}{n\log(2)} \leq \frac{\varepsilon}{\xi\log(2)} \leq \frac{\varepsilon}{5\log(2)} = 0.29\varepsilon, \tag{A.32}$$

since $m \leq n/\xi \leq n/5$. Applying $q \geq nk^2|G|/\varepsilon$ to the second term in Eq. (A.30), we find

$$\frac{(m-2)}{\log(2)} f(k,q) \leq \frac{n}{\xi\log(2)}\left(\frac{\varepsilon}{n} + 2\frac{\varepsilon}{nq} + \frac{\varepsilon}{nq^2}\right) \leq \varepsilon \cdot \frac{1.07}{5\log(2)} = 0.31\varepsilon. \tag{A.33}$$

These error bounds are exactly the same as in Ref. [30] for random unitaries without symmetries. The sum of the two error terms is less than $0.6\varepsilon \leq \varepsilon$ as claimed. □

B Symmetric pseudorandom unitaries

In this section, we present the full details of our results on symmetric pseudorandom unitaries. We begin in Appendix B.1 by reviewing several key definitions and notation. In Appendix B.2, we present our first construction of symmetric pseudorandom unitaries, which has circuit depth $\operatorname{poly} n$ for any discrete on-site symmetry group G. We further prove that our construction can be compiled in terms of constant-local symmetric gates whenever G is Abelian. In Appendix B.3, we combine this result with our gluing theorem from the previous section. This leads to our second construction of symmetric pseudorandom unitaries, which has circuit depth $\operatorname{poly}(\log n)$ on any circuit geometry.

B.1 Definition of symmetric PRUs

We begin by stating the formal definition of a symmetric pseudorandom unitary ensemble. Namely, a symmetric unitary ensemble is a symmetric PRU ensemble if it is indistinguishable from the symmetric Haar ensemble by any bounded-time quantum experiment.

Definition 2 (Symmetric pseudorandom unitaries)**.** *For any symmetry group G acting on n qubits. An infinite sequence $\{\mathcal{U}_n\}_{n\in\mathbb{N}}$ of n-qubit symmetric unitary ensembles $\mathcal{U}_n = \{U_{\mathsf{key}}\}_{\mathsf{key}\in\mathcal{K}_n}$ for the key space $\mathcal{K}_n$ is a symmetric pseudorandom unitary secure against any $t(n)$-time adversary if it satisfies the following.*

- ***Efficient computation:*** *There exists a* $\operatorname{poly}(n)$*-time quantum algorithm that implements the n-qubit unitary U_{key} for all $\mathsf{key} \in \mathcal{K}_n$.*

- ***Indistinguishability from Haar:*** *Any quantum algorithm $\mathcal{A}^U(1^n)$ that runs in time $\leq t(n)$, queries an n-qubit unitary U for any number of times, and outputs $\{0,1\}$ satisfies*

$$\left| \Pr_{\mathsf{key}\sim\mathcal{K}_n} \left[\mathcal{A}^{U_{\mathsf{key}}}(1^n)=1\right] - \Pr_{U\sim H_G} \left[\mathcal{A}^{U}(1^n)=1\right] \right| \leq \mathrm{negl}(n), \tag{B.1}$$

where H_G is the symmetric Haar ensemble and $\mathrm{negl}(n)$ is the negligible function, i.e., a function that is asymptotically smaller than any inverse-polynomial function $1/\mathrm{poly}(n)$.

A $t(n)$-time algorithm $\mathcal{A}^U(1^n)$ is referred to as a $t(n)$-time adversary. The difference between the probability under $\mathcal{U}_n$ and the symmetric Haar ensemble is referred to as the advantage of $\mathcal{A}^U(1^n)$.

This is identical to the definition of a pseudorandom unitary without symmetry when G is trivial.

Our construction of symmetric PRUs will use the notion of a controlled pseudorandom unitary, which are defined as follows. We prove the existence of controlled PRUs in Appendix C.

Definition 3 (Controlled pseudorandom unitaries)**.** *A unitary ensemble is a controlled pseudorandom unitary ensemble secure against any $t(n)$-time adversary if it cannot be distinguished from the controlled Haar ensemble by any $t(n)$-time quantum experiment.*

For brevity, we state the definition colloquially; the full details follow identically to Definition 2. Here, the controlled Haar ensemble refers to a random unitary $U^c = |0\rangle\langle 0| \otimes \mathbb{1} + |1\rangle\langle 1| \otimes U$, where the first qubit is a control register and U is drawn from the Haar measure on the remaining n qubits.

B.2 Proof of Theorem 2: Symmetric PRUs in polynomial depth

We can now present our polynomial-depth construction of symmetric pseudorandom unitaries. Our construction applies to any discrete on-site symmetry group G. Our construction can be compiled in circuit depth $\mathrm{poly}\, n$. This being said, the individual circuit gates in our compilation may not necessarily respect the symmetry operation. For Abelian discrete on-site symmetries in particular, we prove a stronger statement that our construction can be compiled in circuit depth $\mathrm{poly}\, n$ using only symmetric geometrically 5-local gates.

Our construction is simple following the irreducible decomposition of $\mathcal{H}$. We proceed in three steps. First, we apply a unitary circuit that performs the irreducible decomposition, $\mathcal{H} \to \bigoplus_\lambda \mathcal{M}_\lambda \otimes \mathcal{F}_\lambda$. We also compute the irrep index λ onto a small additional register of qubits. Second, conditioned on the small register being in the state λ, we apply a controlled pseudorandom unitary on the register $\mathcal{M}_\lambda$. We repeat this for each irrep λ. Finally, we apply the inverse of the unitary circuit that was applied in the first step.

Our main result is that this construction forms a symmetric pseudorandom unitary.

Theorem 8 (Symmetric pseudorandom unitaries)**.** *For any discrete on-site symmetry group G, the unitary ensemble described above is a symmetric pseudorandom unitary secure against any sub-exponential time quantum adversary.*

We prove the theorem in Appendix B.2.1 below. Our proof relies on several key ingredients whose proofs are provided later sections: the efficient implementation of pseudorandom functions on general local tensor product spaces (Appendix B.2.2), the efficient generation of unitary k-designs on general local tensor product spaces (Appendix B.2.3), and, leveraging these two ingredients, the existence of controlled PRUs (Appendix C). We separate the proofs of these components since they may be of more general interest beyond our current discussion.

For Abelian discrete on-site symmetries, we prove an even stronger statement.

Proposition 2 (Compilation with symmetric local gates)**.** *For any Abelian discrete on-site symmetry G, each random unitary in the ensemble above can be compiled in circuit depth* $\mathrm{poly}(n)$ *using symmetric geometrically 5-local gates.*

Our proof involves a detailed analysis of the unitary circuit in the first step of our symmetric PRU construction, and is provided in Appendix B.2.4. The combination of Theorem 8 and Proposition 2 yields Theorem 2 of the main text.

B.2.1 Proof of Theorem 8: Symmetric PRUs in polynomial depth

Our proof consists of two steps. First, we show that the irreducible decomposition in the first and third steps of our construction can be implemented efficiently. This follows from a relatively straightforward approach, in which we use Clebsch-Gordon coefficients to gradually build up the irreducible decomposition one qudit at a time. Second, we show that controlled PRUs exist and can be implemented in poly n circuit depth on arbitrary local tensor product spaces. This more general setting is necessary because the Hilbert spaces $\mathcal{M}_\lambda$ may not have uniform qudit dimensions.

Let us begin with the first proof step. We work in a 1D circuit geometry for simplicity; our results can be extended to any circuit geometry using Lemma 9 and 10 of Ref. [30]. Our symmetric pseudorandom unitary has the form,

$$U = W^\dagger \cdot V^\dagger \cdot \prod_\lambda U_\lambda^c \cdot V \cdot W, \tag{B.2}$$

Here, we let W denote a unitary circuit that enacts the isomorphism $\mathcal{H} = \bigotimes_{i=1}^n (\mathcal{F}_{\mathrm{reg},i} \otimes \mathcal{A}_i) \cong \mathbb{C}^{D/|G|} \otimes \mathcal{F}_{\mathrm{reg}}$ described in Appendix A.1. Then, we let V denote a circuit that implements the isomorphism $\mathcal{F}_{\mathrm{reg}} \cong \bigoplus_\lambda (\mathbb{C}^{d_\lambda} \otimes \mathcal{F}_\lambda)$ on the output of W. The product VW corresponds to the first step of our symmetric PRU construction.

To implement W efficiently, we recall the behavior of the regular representation of a group under the tensor product. Each qudit Hilbert space has the form $\mathcal{H}_i = \mathcal{F}_{\mathrm{reg}} \otimes \mathcal{A}$ where $|\mathcal{A}| = d'$ and the symmetry group acts in the regular representation on the first register. In Appendix A.1, we reviewed that the tensor product of two regular representations of a group is also regular. This allows one to define an isomorphism,

$$\mathcal{F}_{\mathrm{reg}} \otimes \mathcal{F}_{\mathrm{reg}} \cong \mathbb{C}^{|G|} \otimes \mathcal{F}_{\mathrm{reg}}, \tag{B.3}$$

where the symmetry operator acts only on the second register on the right side. We let W' denote the unitary map that implements this isomorphism. Since W' acts on only a constant-size Hilbert space, it can be implemented in constant depth. To construct the full irreducible decomposition W, we apply W' one-by-one in a ladder-like fashion down the 1D chain. That is, we take $W = W'_{n-1,n} \dots W'_{3,4} W'_{2,3} W'_{1,2}$, where $W'_{i,i+1}$ denotes the application of W' to the $\mathcal{F}_{\mathrm{reg}}$ registers of qudits i and $i+1$. After all $n-1$ applications, the unitary W implements the isomorphism

$$\left(\mathcal{F}_{\mathrm{reg}} \otimes \mathcal{A}\right)^{\otimes n} \cong \left(\mathbb{C}^{|G|} \otimes \mathcal{A}\right)^{\otimes n-1} \otimes \mathcal{F}_{\mathrm{reg}} \otimes \mathcal{A}, \tag{B.4}$$

where the symmetry operator, $W R_g W^\dagger$ where $R_g = \otimes_i R_g^i$, acts only on the $\mathcal{F}_{\mathrm{reg}}$ register of the right side. Since each $W'_{i,i+1}$ can be implemented in constant depth, the unitary W can be implemented in depth $\mathcal{O}(n)$.

To implement the unitary V, we will introduce several ancilla registers. Let $\Lambda = |\{\lambda\}|$ denote the total number of irreps. We introduce two ancilla registers of dimension d_λ for each λ, as well as a final ancilla register of dimension Λ. We let V act on the final $\mathcal{F}_{\mathrm{reg}}$ register and the ancilla registers

as,

$$V\left(|\lambda; x_\lambda, k_\lambda\rangle_{\mathcal{F}_{\text{reg}}} \otimes |0\rangle_\Lambda \otimes \Big(\bigotimes_{\lambda'} |0\rangle_{d_{\lambda'}} \otimes |0\rangle_{\mathcal{F}_{\lambda'}}\Big)\right) \\ = |0\rangle_{\mathcal{F}_{\text{reg}}} \otimes |\lambda\rangle_\Lambda \otimes \big(|x\rangle_{d_\lambda} \otimes |k\rangle_{\mathcal{F}_\lambda}\big) \otimes \Big(\bigotimes_{\lambda' \neq \lambda} |0\rangle_{d_{\lambda'}} \otimes |0\rangle_{\mathcal{F}_{\lambda'}}\Big). \tag{B.5}$$

Here, x_λ labels which multiplicity the state belongs to and k_λ labels the state in the irrep. The unitary V can be implemented in constant depth since it acts on a constant-size Hilbert space. In total, the product VW implements the isomorphism,

$$\big(\mathcal{F}_{\text{reg}} \otimes \mathcal{A}\big)^{\otimes n} \cong \bigoplus_\lambda \Big[\big(\mathbb{C}^{|G|} \otimes \mathcal{A}\big)^{\otimes n-1} \otimes \mathcal{A} \otimes \mathbb{C}^{d_\lambda}\Big] \otimes \mathcal{F}_\lambda, \tag{B.6}$$

where we identify the Hilbert space in square brackets as $\mathcal{M}_\lambda$. Note that $\mathcal{M}_\lambda$ has a local tensor product structure, but with a non-uniform set of local Hilbert space dimensions $|G|, d', d_\lambda$.

It remains only to show that the second step of our symmetric pseudorandom unitary construction can be implemented efficiently. In Theorem 11 of Appendix C, we show that controlled PRUs exist on any Hilbert space that factorizes as $\mathcal{H} = \mathcal{H}_L \otimes \mathcal{H}_R$. Here, we have in mind $\mathcal{H} \to \mathcal{M}_\lambda$ for each λ. The security of our controlled PRU is determined by the dimensions of $\mathcal{H}_L$ and $\mathcal{H}_R$; in particular, when both spaces have dimension exponential in n, we achieve security against any sub-exponential time quantum adversary. Our construction assumes only two ingredients: an efficient implementation of a quantum-secure pseudorandom function on $\mathcal{H}_L$ and $\mathcal{H}_R$, and an efficient implementation of an approximate unitary 2-design on $\mathcal{H}$. In Appendix B.2.2, we show that existing results allow one to generate quantum-secure pseudorandom functions in $\operatorname{poly} n$ depth on any local tensor product space. In Appendix B.2.3, we introduce a new, nested application of the two-layer circuit construction from Ref. [30] which allows one to generate unitary k-designs in $\operatorname{poly}(\log n)$ depth on any local tensor product space. Since we have shown that $\mathcal{M}_\lambda$ is a local tensor product space, these results prove that controlled PRUs with sub-exponential security can be implemented in $\operatorname{poly} n$ depth on each $\mathcal{M}_\lambda$. This completes our proof. □

B.2.2 Pseudorandom functions on any local tensor product space

In this brief section, we show that existing results imply that the first three components, F, O_L, and O_R, of our controlled PRU construction can be realized in $\operatorname{poly} n$ depth in any local tensor product space. This extends known results for systems composed solely of qubits [30, 34]. Here, a local tensor product space is given by $\mathcal{H} = \bigotimes_{i=1}^n \mathbb{C}^{q_i}$, for arbitrary constant-sized local qudit dimensions q_i.

Fact 1. *One can generate the operator F in circuit depth $\operatorname{poly} n$ on any local tensor product space and any circuit geometry.*

Proof. Ref. [113] constructs a quantum-secure pseudorandom function $f : [D] \to \{0,1\}$ for $D = 2^\ell$. The function can be compiled in circuit depth $\operatorname{poly} n$ on any circuit geometry [30]. Setting $\ell = \lceil D \rceil$ and restricting the action of the function to the first D bits yields F. □

Fact 2. *One can generate the operators O_L, O_R in circuit depth $\operatorname{poly} n$ on any local tensor product space and any circuit geometry.*

Proof. Ref. [113] constructs a quantum-secure pseudorandom function f_L for $D_L = 2^\ell$ and $D_R = p$, where p is any $\operatorname{poly} n$-bit integer. The function can be compiled in circuit depth $\operatorname{poly} n$ on any circuit geometry [30]. Setting $\ell = \lceil D_L \rceil$ and restricting the action of the function to the first D_L bits yields O_L, and similar for O_R. □

B.2.3 Unitary designs on any local tensor product space

In this section, we show that unitary k-designs can be generated in $\mathrm{poly}(\log n)$ depth on any local tensor product space. To do so, we introduce a *nested* variant of the two-layer construction from Ref. [30], in which the small random unitaries of the two-layer circuit are themselves decomposed into smaller two-layer circuits. This enables a stand-alone proof that unitary designs can be realized in extremely low depth, which does not rely on any existing results on the convergence of local random circuits to designs. This is convenient for our purposes, since it avoids the need to perform a new spectral gap analysis for general local tensor product spaces.

Our main result is the following.

Proposition 3 (Unitary designs on any local tensor product space)**.** *One can generate ε-approximate unitary k-designs in circuit depth $\mathcal{O}(\log(n)\cdot\log(1/\varepsilon)\cdot\log(nk/\varepsilon)^4\cdot k^8)$ on any local tensor product space and any circuit geometry.*

Proof. We arrange the n qudits into a 1D line. We then divide the qubits into n/ξ contiguous patches, so that each patch has Hilbert space dimension at least 2^ξ for a chosen ξ. We then sub-divide each patch into ξ/ζ contiguous smaller patches, so that each smaller patch has Hilbert space dimension at least 2^ζ for some $\zeta < \xi < n$. We consider a two-layer circuit on the large patches, repeated β times. Each small random unitary of the two-layer circuit is generated, in turn, by a repeated two-layer circuit on the smaller patches, repeated α times. Each small random unitary in the small two-layer circuits is taken to be Haar-random. The construction has circuit depth $\mathcal{O}(\alpha\cdot\beta\cdot 2^{4\zeta})$, since a Haar-random unitary on a Hilbert space of dimension $\mathcal{O}(2^{2\zeta})$ can be compiled in circuit depth $\mathcal{O}(2^{4\zeta})$.

To determine the required values of α, β, ξ, ζ, let us begin by analyzing the smaller repeated two-layer circuits. From Theorem 1 of Ref. [30], each small two-layer circuit within each patch forms a unitary k-design with relative error $\varepsilon_0' \le \xi k^2/2^\zeta$. Let us set $\zeta \ge \log_2(3\xi k^2)$ so that $\varepsilon_0' \le 1/3$. From Fact 3 below, this implies that the repeated small two-layer circuit within each patch has relative error less than $2^{\alpha-1}\varepsilon_0'^\alpha \le (2/3)^\alpha$. We can then set $\alpha \ge \log_{3/2}(1/\varepsilon')$ to obtain a ε'-approximate unitary k-design on each patch.

We can now apply Theorem 1 of Ref. [30] again, to analyze the larger two-layer circuits. To do so, we set $\varepsilon' \le \varepsilon_0/n$ and $\xi \ge \log_2(nk^2/\varepsilon)$. This guarantees that the nested two-layer circuit is an ε_0-approximate unitary k-design on all n qubits. Similar to the previous paragraph, let us set $\varepsilon_0 \le 1/3$. From Fact 3, the repeated large two-layer circuit has relative error less than $2^{\beta-1}\varepsilon_0^\beta \le (2/3)^\beta$. We can then set $\beta \ge \log_{3/2}(1/\varepsilon)$ to obtain a ε-approximate unitary k-design on all n qubits. In total, our design has circuit depth $\mathcal{O}(\alpha\cdot\beta\cdot 2^{4\zeta}) = \mathcal{O}(\log(1/\varepsilon')\cdot\log(1/\varepsilon)\cdot(3\xi k^2)^4) = \mathcal{O}(\log(n)\cdot\log(1/\varepsilon)\cdot\log(nk/\varepsilon)^4\cdot k^8)$. This completes the proof. □

Our proof above uses the following fact. The fact allows one to exponentially suppress the relative error of any unitary design by repeatedly applying the design several times in a row.

Fact 3. *If $\Phi_\mathcal{E}$ has relative error ε, then $\Phi_\mathcal{E}\circ\Phi_\mathcal{E}$ has relative error less than $2\varepsilon^2$.*

Proof. We write $\Phi_\mathcal{E} = \Phi_H + \delta\Phi$, where $\Phi_H\circ\delta\Phi = \delta\Phi\circ\Phi_H = 0$ by definition. Squaring gives, $\Phi_\mathcal{E}\circ\Phi_\mathcal{E} = \Phi_H + \delta\Phi\circ\delta\Phi$. By assumption, we have $-\varepsilon\Phi_H(\rho) \preceq \delta\Phi(\rho) \preceq \varepsilon\Phi_H(\rho)$ for any ρ. We can break $\delta\Phi(\rho) = \rho_+ - \rho_-$ into a positive and negative part, $0 \preceq \rho_\pm \preceq \varepsilon\Phi_H(\rho)$. We then have $\delta\Phi(\delta\Phi(\rho)) = \delta\Phi(\rho_+) - \delta\Phi(\rho_-)$. Moreover, since $\Phi_\mathcal{E}$ has relative error ε, we can bound $\delta\Phi(\rho_\pm) \preceq \varepsilon\Phi_H(\rho_\pm)$. Combining these two inequalities, we have

$$-2\varepsilon^2\Phi_H(\rho) \preceq -\varepsilon\Phi_H(\rho_+) - \varepsilon\Phi_H(\rho_-) \preceq \delta\Phi(\delta\Phi(\rho)) \preceq \varepsilon\Phi_H(\rho_+) + \varepsilon\Phi_H(\rho_-) \preceq 2\varepsilon^2\Phi_H(\rho), \tag{B.7}$$

which completes the proof. □

B.2.4 Proof of Proposition 2: Compiling with symmetric geometrically 5-local gates

Let us begin by recalling a few facts about discrete on-site Abelian symmetries. Any finite Abelian group is isomorphic to a tensor sum of cyclic groups,

$$G \cong \mathbb{Z}_{p_1} \otimes \mathbb{Z}_{p_2} \otimes \cdots \otimes \mathbb{Z}_{p_l}, \tag{B.8}$$

for positive numbers $p_1, \ldots, p_l$. Each irreducible representation x of an Abelian group has dimension $d_x = 1$. This implies that each irrep x appears once in the regular representation, so we can write $\mathcal{F}_{\text{reg}} = \{|x\rangle\}$. It also implies that for any fixed irrep, the action of any group element amounts to an overall phase. The label x of each irrep is referred to as its *charge*,

$$x \in \mathbb{Z}_{p_1} \otimes \mathbb{Z}_{p_2} \otimes \cdots \otimes \mathbb{Z}_{p_l}, \tag{B.9}$$

and determines the overall phase accrued when a symmetry group element is applied, $R_g |x\rangle = \prod_{j=1}^{l} e^{i(2\pi/p_j)x_j g_j} |x\rangle$. From this definition, it is clear that the charges of irreps add under the tensor product, i.e. the state $|x_1\rangle \otimes |x_2\rangle$ has charge $x_1 \oplus x_2$, where $\oplus$ denotes addition in the group G.

Let us now turn to compilation of our symmetric random unitary ensemble. We note two facts. First, the unitary V is equal to the identity for an Abelian group, since the regular representation is already a tensor sum with each irrep appearing once. Second, we can take the unitary W to be a generalized CNOT ladder. Namely, we let $W'_{i,i+1}(|y_i\rangle \otimes |x_{i+1}\rangle) = |y_i\rangle \otimes |x_{i+1} \oplus y_i\rangle$ for each i. After n applications, one finds $W |x_1 \ldots x_n\rangle = |y_1 \ldots y_n\rangle$, where $y_i = x_1 \oplus x_2 \oplus \cdots \oplus x_i$ is equal to the charge of all qudits to the left of i (including i itself). The n-th qudit contains the charge y_n of the entire system, as required by the definition of W.

To implement our symmetric unitary ensemble, $U = W^\dagger(\prod_{y_n} U^c_{y_n})W$, we first decompose each controlled PRU, $U^c_{y_n}$, as a product of controlled unitary gates of the form,

$$|y_n\rangle\langle y_n| \otimes V_{i,i+1} + \big(\mathbb{1} - |y_n\rangle\langle y_n|\big) \otimes \mathbb{1}_{i,i+1}. \tag{B.10}$$

Here, the projector $|y_n\rangle\langle y_n|$ acts on the n-th qudit (or any register that stores the total charge y_n), and $V_{i,i+1}$ acts on qudits $i, i+1$. We can assume that each $V_{i,i+1}$ is a two-qudit geometrically-local gate since such gates are universal for quantum computation.

To proceed, we consider the conjugation of the gate $V_{i,i+1}$ by the final $n-i+1$ layers of the circuit W. That is,

$$\tilde{V}_{i,i+1} \equiv W^\dagger_{i-1,i} \ldots W^\dagger_{n-1,n} V_{i,i+1} W_{n-1,n} \ldots W_{i-1,i}. \tag{B.11}$$

We let y_j denote the value of qudit j after the W gates are applied (i.e. where $V_{i,i+1}$ acts), and x_j the value before the W gates are applied. We note three properties. First, the light-cone of the W gates implies that $\tilde{V}_{i,i+1}$ can act at most on qudits $\{i-1, i, i+1, i+2\}$. Second, the qudit $\{i-1\}$ must function only as a control qudit for an operation on the remaining $\{i, i+1, i+2\}$. That is, we can write $\tilde{V}_{i,i+1} = \sum_{y_i} |y_{i-1}\rangle\langle y_{i-1}| \otimes \tilde{V}^{y_i}_{i,i+1,i+2}$. This follows because the value of y_{i-1} is not changed by the gate $W_{i-1,i}$ nor the unitary $V_{i,i+1}$ (since $V_{i,i+1}$ acts only on y_i and y_{i+1} and not y_{i-1}). Third, the controlled operations $\tilde{V}^{y_i}_{i,i+1,i+2}$ are symmetric on $\{i, i+1, i+2\}$. This follows because $V_{i,i+1}$ cannot change the value of $y_{i+2} = y_{i-1} \oplus x_i \oplus x_{i+1} \oplus x_{i+2}$. Since $V_{i,i+1}$ also cannot change the value of y_{i-1}, it must not change the charge, $x_i \oplus x_{i+1} \oplus x_{i+2}$, on $\{i, i+1, i+2\}$. This implies that $\tilde{V}^{y_i}_{i,i+1,i+2}$ is symmetric.

With these three properties established, our compilation of U is as follows. We copy the charge y_n of all qudits to a first ancilla register. We then iterate over all charges y_n and all gates in $U^c_{y_n}$ [Eq. (B.10)]. For each gate, we first copy the charge y_{i-1} to a second ancilla register. We then perform the gate $\tilde{V}_{i,i+1}$ as a controlled gate from the first and second ancilla register to qudits $\{i, i+1, i+2\}$. From the discussion above, this is a symmetric gate acting on three geometrically-local system qudits

and two ancilla qudits. We then uncompute the charge y_{i-1} and repeat for all charges y_n and all gates in $U^c_{y_n}$. After all gates are completed, we uncompute the charge y_n. This completes our compilation of U. Our compilation has a factor of n greater depth compared to the original compilation, due to the need to compute y_{i-1} for every gate. Hence, if the original circuit depth is $\operatorname{poly} n$, our symmetric compilation also has circuit depth $n \cdot \operatorname{poly} n = \operatorname{poly} n$.

In the above discussion, we neglected to notate the ancilla register $\mathcal{A}$ to reduce notation. One can easily see that the analysis is unchanged by the addition of new registers that are not acted on by the symmetry group, since these registers are also not acted on by any of the $W_{i,i+1}$. □

B.3 Proof of Theorem 3: Symmetric PRUs in poly-logarithmic depth

The combination of our gluing theorem (Theorem 7) and our symmetric pseudorandom unitary construction (Theorem 8) immediately implies that symmetric pseudorandom unitaries can be generated in extremely low circuit depths. This is formalized in Theorem 9 and Corollary 1 below, which together prove Theorem 3 of the main text. Our proofs of both statements are extremely short building upon our earlier results.

Theorem 9 (Gluing small symmetric pseudorandom unitaries)**.** *For any discrete on-site symmetry group G. Let n be the number of qudits in the whole system and $\xi = \omega(\log n)$ be the number of qudits in each local patch. Suppose each small random unitary in the two-layer brickwork ensemble $\mathcal{E}$ is a 2ξ-qubit symmetric PRU secure against $\operatorname{poly}(n)$-time adversaries. Then $\mathcal{E}$ is an n-qubit symmetric PRU secure against $\operatorname{poly}(n)$-time adversaries.*

Proof. Our proof is identical to the proof of Theorem 2 in Ref. [30]. The only change is to replace Theorem 1 in Ref. [30], which applies to random unitaries without symmetries, with Theorem 7 in our work, which applies to random unitaries with symmetries. □

Corollary 1 (Low-depth pseudorandom unitaries)**.** *For any discrete on-site symmetry group G. Under the conjecture that no subexponential-time quantum algorithm can solve LWE, random quantum circuits over n qubits can form symmetric PRUs secure against any polynomial-time quantum adversary in circuit depth $d = \operatorname{poly} \log n$, on any circuit geometry.*

Proof. The corollary follows immediately from Theorem 8 and Theorem 9. Here, we take the number of qudits in Theorem 8 to be $\xi = \operatorname{poly}(\log n)$. □

C Controlled pseudorandom unitaries

In this section, we take a brief detour from our analysis of symmetric systems to prove the existence of controlled PRUs. Controlled PRUs are unitary ensembles that are indistiguishable from controlled Haar-random unitaries by any polynomial-time quantum adversary. They are also central to our construction of symmetric PRUs in the previous section. Nonetheless, to this point, the existence of a controlled PRU ensemble is not known in the literature. In what follows, we analyze controlled variants of the PFC and LRFC PRU ensembles, and prove that both controlled ensembles are controlled PRUs. Our analysis requires multiple technical innovations compared to existing PRU proofs to handle the control register.

C.1 Summary of results

Let us first define several operators and notations. We consider a Hilbert space $\mathcal{H} = \mathcal{H}_{\mathsf{C}} \otimes \mathcal{H}_{\mathsf{S}}$ where $\mathcal{H}_{\mathsf{C}}$ is a single-qubit control register and $\mathcal{H}_{\mathsf{S}}$ is a system register of dimension D. The Hilbert

space is spanned by the computational basis states, $\{|t, x\rangle \mid t \in \{0,1\}, x \in [D]\}$. We let F^c denote a quantum-secure controlled pseudorandom function (PRF) acting on this Hilbert space in the phase representation,

$$F^c |0, x\rangle = |0, x\rangle , \qquad F^c |1, x\rangle = (-1)^{f(x)} |1, x\rangle , \tag{C.1}$$

where $f : [D] \to \{0, 1\}$. We let P^c denote a quantum-secure controlled pseudorandom permutation (PRP),

$$P^c |0, x\rangle = |0, x\rangle , \qquad P^c |1, x\rangle = |1, P(x)\rangle , \tag{C.2}$$

where $P : [D] \to [D]$ is one-to-one. For any Hilbert space that factorizes, $\mathcal{H} = \mathcal{H}_L \otimes \mathcal{H}_R$, we let S^c_L denote quantum-secure controlled PRFs, acting in the representation,

$$S^c_L |0, x_L, x_R\rangle = |0, x_L, x_R\rangle , \qquad S^c_L |1, x_L, x_R\rangle = |1, x_L \oplus f_L(x_R), x_R\rangle , \tag{C.3}$$

where $f_L : [D_R] \to [D_L]$. We let S^c_R denote the analogous PRF,

$$S^c_R |0, x_L, x_R\rangle = |0, x_L, x_R\rangle , \qquad S^c_R |1, x_L, x_R\rangle = |1, x_L, x_R \oplus f_R(x_L)\rangle , \tag{C.4}$$

where $f_R : [D_L] \to [D_R]$. Here, $D_L = |\mathcal{H}_L|$, $D_R = |\mathcal{H}_R|$, and $D_L D_R = D = |\mathcal{H}|$. We note that quantum-secure controlled PRFs and PRPs are easily constructed from quantum-secure PRFs and PRPs without controls by using controlled SWAP gates[3]. Finally, we let C^c denote a controlled application of any unitary 2-design,

$$C^c = |0\rangle\langle 0| \otimes \mathbb{1} + |1\rangle\langle 1| \otimes C, \tag{C.5}$$

where $C \sim \mathfrak{D}$ and $\mathfrak{D}$ is a 2-design.

Definitions in hand, we consider the following two controlled random unitary ensembles. First, we consider a controlled variant of the Permutation-Function-Clifford (PFC) ensemble [32, 34],

$$U^c = P^c \cdot F^c \cdot C^c, \tag{C.6}$$

where P^c, F^c, C^c are drawn randomly as described above. The PFC ensemble without controls is the simplest known ensemble that forms a PRU. Second, we consider a controlled variant of the Luby-Rackoff-Function-Clifford (LRFC) ensemble [35],

$$U^c = S^c_R \cdot F^c \cdot S^c_L \cdot C^c. \tag{C.7}$$

[3]In more detail, to obtain a controlled PRF F^c from a standard PRF F, we consider the operation cSWAP$(\mathbb{1} \otimes \mathbb{1} \otimes F)$cSWAP, where the three subsystems are the control qubit, the system register, and an ancilla register of the same size as the system, respectively. Here, cSWAP performs a controlled swap between the system and ancilla registers. If one prepares the ancilla register in the zero state, this operator is equal to the desired controlled PRF up to a controlled phase $(-1)^{f(0)}$. This phase can be determined and undone with one additional query to f.

To obtain a controlled PRP P^c from a standard PRP, we consider the operation cSWAP$(\mathbb{1} \otimes \mathbb{1} \otimes P)$cSWAP, and assume that the ancilla register again begins in the zero state. If the control register is in the 1 state, this is equal to the desired controlled PRP. If the control register is in the 0 state, it is equal to the desired PRP on the control and system registers, but leaves the ancilla register in the state $|P(0)\rangle$. The ancilla register can then be traced out, leaving a controlled PRP on the control and system. Alternatively, to maintain unitarity on the entire system including the ancilla qubits, one can apply another query to P on an additional ancilla register prepared in the all zero state, use the result to controllably return the original ancilla register to the zero state, and then apply P^{-1} to the additional ancilla register to return that register to the zero state as well.

The controlled PRFs O^c_L and O^c_R can be obtained from their standard versions in an identical fashion to controlled PRPs.

The LRFC ensemble was introduced as an alternative to the PFC ensemble, which relies solely on pseudorandom functions and not permutations. This can be convenient, as existing results on quantum-secure PRPs are more limited than those on PRFs. In particular, for our application to symmetric pseudorandom unitaries in this work, we require the existence of controlled PRUs on systems composed of qudits of arbitrary dimension. Existing analyses of quantum-secure PRPs are limited to qu*bit* systems; hence, our main results will utilize only on the LRFC ensemble.

Our main result is that both of these controlled random unitary ensembles are controlled PRUs.

Theorem 10. *The controlled PFC ensemble is a controlled PRU ensemble, with security against any sub-exponential time quantum adversary.*

Theorem 11. *The controlled LRFC ensemble is also a controlled PRU ensemble, with security against any sub-exponential time quantum adversary.*

We prove the theorems in the following two subsections. Both of our proofs follow the concise strategy introduced in Ref. [44] for proving the adaptive security of PRUs. We augment this approach with several technical innovations to extend the analysis to controlled random unitaries. This extension must be performed carefully, since the control register can take exponentially many values, $|\{0,1\}^k| = 2^k$, in an experiment that queries the unitary $k = \operatorname{poly} n$ times. Thus, many naive proof approaches will diverge exponentially in k. We solve this by keeping a detailed accounting of the controlled register and leveraging several useful operator inequalities. Our proof of Theorem 10 is simpler and is presented first. Our proof of Theorem 11 then follows by similar steps.

C.2 Proof of Theorem 10: Controlled-PFC ensemble

We can express the output state of any experiment that queries U_c up to k times as,

$$|\psi_U\rangle_{\mathsf{AB}} = (2D)^k \cdot (\mathbb{1}_{\mathsf{AB}} \otimes \langle\Psi_{\text{Bell}}|_{\mathsf{XY}})(\mathbb{1}_{\mathsf{ABY}} \otimes \boldsymbol{U}_{\mathsf{X}}) |\Psi\rangle_{\mathsf{ABXY}}\,, \tag{C.8}$$

where $\boldsymbol{U} \equiv (U^c)^{\otimes k}$. We also denote $\boldsymbol{P} = (P^c)^{\otimes k}$, $\boldsymbol{F} = (F^c)^{\otimes k}$, $\boldsymbol{C} = (C^c)^{\otimes k}$. We refer to Ref. [30] for a visual depiction. Here, A is the combined control and system register, B is an arbitrary ancilla register, and X and Y are ancilla registers of the same size as k copies of A. $|\Psi_{\text{Bell}}\rangle_{\mathsf{XY}}$ is a Bell state between X and Y. The factor $(2D)^k = |\mathsf{X}| = |\mathsf{Y}|$ normalizes the state after the Bell projection.

To write the expected output state in a more compact form, we let

$$B'_{\mathsf{XY}} \equiv (2D)^{2k} \cdot \mathop{\mathbb{E}}_{P,f} \left[(\boldsymbol{F}^\dagger \boldsymbol{P}^\dagger \otimes \mathbb{1})_{\mathsf{XY}} |\Psi_{\text{Bell}}\rangle\langle\Psi_{\text{Bell}}|_{\mathsf{XY}} (\boldsymbol{P}\boldsymbol{F} \otimes \mathbb{1})_{\mathsf{XY}} \right] \tag{C.9}$$

denote the Bell projector twirled over PF, multiplied by $(2D)^{2k}$. This yields,

$$\rho_{\mathsf{AB}} \equiv \mathop{\mathbb{E}}_{P,f} \left[|\psi_U\rangle\langle\psi_U|_{\mathsf{AB}} \right] = \mathop{\mathbb{E}}_{C} \left[\operatorname{tr}_{\mathsf{XY}} \left(B'_{\mathsf{XY}} \boldsymbol{C}_{\mathsf{X}} |\Psi\rangle\langle\Psi|_{\mathsf{ABXY}} \boldsymbol{C}_{\mathsf{X}}^\dagger \right) \right]. \tag{C.10}$$

We let $\rho^H_{\mathsf{AB}} \equiv \mathbb{E}_{U\sim H} |\psi_U\rangle\langle\psi_U|_{\mathsf{AB}}$ denote the output of the same experiment when U_c is controlled Haar-random. We will show that $\|\rho - \rho^H\|_1 \leq 10\delta$, where $\delta = k^2/2^n$. This establishes security against any sub-exponential time quantum adversary.

We begin by writing down several definitions and facts. We define the controlled distinct subspace projectors,

$$\Pi^{\text{c-dist}}_{\mathsf{X}} = \sum_{t\in\{0,1\}^k} |t\rangle\langle t|_{\mathsf{X_C}} \otimes \left[\Pi^{\text{dist}}_t \otimes \mathbb{1}_{\bar{t}}\right]_{\mathsf{X_S}}, \qquad \Pi^{\text{c-dist}}_{\mathsf{Y}} = \sum_{t\in\{0,1\}^k} |t\rangle\langle t|_{\mathsf{Y_C}} \otimes \left[\Pi^{\text{dist}}_t \otimes \mathbb{1}_{\bar{t}}\right]_{\mathsf{Y_S}}, \tag{C.11}$$

as well as their product,

$$\Pi^{\text{c-dist}}_{\mathsf{XY}} = \Pi^{\text{c-dist}}_{\mathsf{X}} \otimes \Pi^{\text{c-dist}}_{\mathsf{Y}}. \tag{C.12}$$

Here, $\mathsf{X}_{\mathsf{L},\mathsf{C}}$ denotes the k control registers in X_{L}, and $\mathsf{X}_{\mathsf{L},\mathsf{S}}$ the k system register, and similar for Y_{R}. Here, $\Pi_t^{\text{dist}} = \sum_{x\in\text{dist}} |x\rangle\langle x|$ is the projector onto the distinct subspace on the system registers of copies j with $t_j = 1$, where sum runs over all distinct bitstrings in $[D_L]^{\otimes|t|}$ or $[D_R]^{\otimes|t|}$, and $\mathbb{1}_{\bar{t}}$ is the identity matrix on the copies j with $t_j = 0$.

In analogy to B' above, we also define the operator,

$$B_{\mathsf{XY}} = D^{2(k-|t|)} \sum_{\substack{t,t'\in\{0,1\}^k \\ |t|=|t'|}} \left(|t',t'\rangle\langle t,t|_{\mathsf{X_C Y_C}} \otimes \sum_{\pi\in S_{|t|}} \left[(\pi_{t\to t'} \otimes \pi_{t\to t'}) \otimes |\Psi_{\text{Bell}}^{\bar{t}'}\rangle\langle\Psi_{\text{Bell}}^{\bar{t}}| \right]_{\mathsf{X_S Y_S}} \right), \quad \text{(C.13)}$$

where $\pi_{t\to t'}$ is a permutation mapping the system registers of the copies of X where $t_j = 1$ to the those where $t'_j = 1$, and similarly for Y. With these definitions in hand, our proof uses five facts:

(1) B and B' are equal on the controlled distinct subspace up to small relative error, $B_{\mathsf{XY}}\Pi_{\mathsf{XY}}^{\text{c-dist}} \preceq B'_{\mathsf{XY}}\Pi_{\mathsf{XY}}^{\text{c-dist}} \preceq (1+\delta)B_{\mathsf{XY}}\Pi_{\mathsf{XY}}^{\text{c-dist}}$. To see this, we express the former as

$$B'_{\mathsf{XY}} = \sum_{\substack{t,t'\in\{0,1\}^k \\ |t|=|t'|}} D^{2(k-|t|)} \cdot |t,t\rangle\langle t',t'|_{\mathsf{X_C Y_C}} \otimes \left[(\tilde{\pi}_{t\to t'} \otimes \tilde{\pi}_{t\to t'}) \cdot \left(B'_t \otimes |\Psi_{\text{Bell}}^{\bar{t}}\rangle\langle\Psi_{\text{Bell}}^{\bar{t}}| \right) \right]_{\mathsf{X_S Y_S}}, \quad \text{(C.14)}$$

where $\tilde{\pi}_{t\to t'}$ is any permutation that maps copies j where $t_j = 1$ to those where $t'_j = 1$. We also abbreviate,

$$(B'_t)_{\mathsf{XY}_t} \equiv D^{2|t|} \mathop{\mathbb{E}}_{P,f} \left[(F^{\dagger,\otimes|t|} P^{\dagger,\otimes|t|} \otimes \mathbb{1})_{\mathsf{XY}_t} |\Psi_{\text{Bell}}^t\rangle\langle\Psi_{\text{Bell}}^t|_{\mathsf{XY}_t} (P^{\otimes|t|} F^{\otimes|t|} \otimes \mathbb{1})_{\mathsf{XY}_t} \right]. \quad \text{(C.15)}$$

We can perform precisely the same decomposition for B, replacing $(B'_t)_{\mathsf{XY}_t}$ with $(B_t)_{\mathsf{XY}_t} \equiv \sum_{\pi\in S_{|t|}} \pi\otimes\pi$. Inserting the controlled distinct subspace projectors, we find

$$\begin{aligned}
B'_t(\Pi_{\mathsf{X}_t}^{\text{dist}} \otimes \Pi_{\mathsf{Y}_t}^{\text{dist}}) &= D^{|t|} \mathop{\mathbb{E}}_{P,f} \sum_{x\,\text{dist}} \sum_{\tilde{x}} (-1)^{f(x)+f(\tilde{x})} |\tilde{x}, P\tilde{x}\rangle\langle x, Px|_{\mathsf{XY}_t} \\
&= D^{|t|} \mathop{\mathbb{E}}_{P} \sum_{\pi} \sum_{x\,\text{dist}} |\pi x, \pi Px\rangle\langle x, Px|_{\mathsf{XY}_t} \\
&= (D^{|t|}/\mathfrak{D}_{|t|}) \sum_{\pi} \sum_{x\,\text{dist}} \sum_{y\,\text{dist}} |\pi x, \pi y\rangle\langle x, y|_{\mathsf{XY}_t} \\
&= (D^{|t|}/\mathfrak{D}_{|t|}) \sum_{\pi} (\pi\otimes\pi)_{\mathsf{XY}_t} \Pi_{\mathsf{XY}_t}^{\text{c-dist}} = (D^{|t|}/\mathfrak{D}_{|t|})(B_t)_{\mathsf{XY}_t} (\Pi_{\mathsf{X}_t}^{\text{dist}} \otimes \Pi_{\mathsf{Y}_t}^{\text{dist}}),
\end{aligned}$$

where the expectation over f enforces $\tilde{x} = \pi x$ for some $\pi \in S_{|t|}$, and the expectation over P yields a random distinct bitstring y. In the final line, we use $\sum_{x,y} |x,y\rangle\langle x,y| = \mathbb{1}$ to eliminate the sums. We see that both B'_t and B_t are block diagonal in $|t|$, and are equal in each block up to the constant factor $D^{|t|}/\mathfrak{D}_{|t|}$. The inequality $1 \le D^{|t|}/\mathfrak{D}_{|t|} \le 1+\delta$ completes our claim.

(2.) The twirl of $\Pi_{\mathsf{X}}^{\text{c-dist}}$ over a controlled 2-design is close to the identity, $\left\| \mathbb{1} - \mathbb{E}_{C\sim\mathfrak{D}} \left[\boldsymbol{C}_{\mathsf{X}}{}^{\dagger} \Pi_{\mathsf{X}}^{\text{c-dist}} \boldsymbol{C}_{\mathsf{X}} \right] \right\|_\infty \le \delta$. To show this, we bound $\mathbb{1} - \Pi_{\mathsf{X}}^{\text{c-dist}} \le \sum_{t\in\{0,1\}^k} \sum_{i<j\in t} |t\rangle\langle t|_{\mathsf{X_C}} \otimes \Pi_{\mathsf{X_S}}^{\mathsf{eq},ij}$, where

$$\Pi_{\mathsf{X_S}}^{\mathsf{eq},ij} = \sum_x |x,x\rangle\langle x,x|_{\mathsf{X}_{\mathsf{S},i}\mathsf{X}_{\mathsf{S},j}}.$$

The twirl over a controlled 2-design yields $\mathbb{E}_C[\boldsymbol{C}_{\mathsf{X}}{}^{\dagger} \sum_t \sum_{i<j\in t} |t\rangle\langle t| \otimes \Pi_{\mathsf{X_S}}^{\mathsf{eq},ij} \boldsymbol{C}_{\mathsf{X}}] = \sum_t \sum_{i<j\in t} |t\rangle\langle t| \otimes (\mathbb{1} + \mathcal{S}_{ij})/(D+1)$, where $\mathcal{S}_{ij}$ is the swap operator. Hence, the spectral norm is less $\max_t \sum_{i<j\in t} 2/(D+1) \le (k^2/2)(2/D) = \delta$.

(3.) B and B' commute with $\Pi_{\mathsf{X}}^{\text{c-dist}}$ and $\Pi_{\mathsf{Y}}^{\text{c-dist}}$. B also commutes with $\boldsymbol{C}$.

(4.) For any positive Hermitian matrices P, Q, R, where P and Q commute, we have $\operatorname{tr}(PQR) \leq \|Q\|_\infty \operatorname{tr}(PR)$. This follows from Holder's inequality,

$$\operatorname{tr}(PQR) = \operatorname{tr}(Q\sqrt{P}R\sqrt{P}) \leq \|Q\|_\infty \|\sqrt{P}R\sqrt{P}\|_1 = \|Q\|_\infty \operatorname{tr}(PR).$$

(5.) We have $1 - \delta \leq \operatorname{tr}\big(B_{\mathsf{XY}} \Pi_{\mathsf{Y}}^{\text{c-dist}} |\Psi\rangle\langle\Psi|\big) \leq 1$. This is shown below the current proof.

We can now prove the claim. We insert $\mathbb{1} = \Pi_{\mathsf{XY}}^{\text{c-dist}} + (\mathbb{1} - \Pi_{\mathsf{XY}}^{\text{c-dist}})$ to decompose ρ into two terms,

$$\rho = \rho_{\text{dist}} + \delta\rho \equiv \mathop{\mathbb{E}}_{C} \big[\operatorname{tr}_{\mathsf{XY}}(B' \cdot \Pi_{\mathsf{XY}}^{\text{c-dist}} \cdot \boldsymbol{C} |\Psi\rangle\langle\Psi| \boldsymbol{C}^\dagger)\big] + \mathop{\mathbb{E}}_{C} \big[\operatorname{tr}_{\mathsf{XY}}(B' \cdot (\mathbb{1} - \Pi_{\mathsf{XY}}^{\text{c-dist}}) \cdot \boldsymbol{C} |\Psi\rangle\langle\Psi| \boldsymbol{C}^\dagger)\big].$$

We bound the second term as follows,

$$\begin{aligned}
\|\delta\rho\|_1 = \operatorname{tr}(\delta\rho) &= \mathop{\mathbb{E}}_{C} \big[\operatorname{tr}\Big(B'(\mathbb{1} - \Pi_{\mathsf{XY}}^{\text{c-dist}})\boldsymbol{C} |\Psi\rangle\langle\Psi| \boldsymbol{C}^\dagger\Big)\big] && \text{(since } \delta\rho \text{ is positive)} \\
&\leq 1 - \mathop{\mathbb{E}}_{C} \big[\operatorname{tr}\big(B \cdot \Pi_{\mathsf{XY}}^{\text{c-dist}} \cdot \boldsymbol{C} |\Psi\rangle\langle\Psi| \boldsymbol{C}^\dagger\big)\big], && \text{(from } \operatorname{tr}(\rho) = 1 \text{ and (1.))} \\
&= 1 - \operatorname{tr}\big(B \cdot \mathop{\mathbb{E}}_{C} \big[\boldsymbol{C}^\dagger \Pi_{\mathsf{X}}^{\text{c-dist}} \boldsymbol{C}\big] \otimes \Pi_{\mathsf{Y}}^{\text{c-dist}} \cdot |\Psi\rangle\langle\Psi|\big), && \text{(from (3.))} \\
&\leq 1 - (1 - \delta) \operatorname{tr}\big(B(\mathbb{1} \otimes \Pi_{\mathsf{Y}}^{\text{c-dist}}) |\Psi\rangle\langle\Psi|\big) && \text{(from (2.), (3.), (4.))} \\
&= 1 - (1 - \delta)^2 \leq 2\delta. && \text{(from (5.))}
\end{aligned}$$

Meanwhile, the first term is close to a fixed density matrix ρ_a, independent of the ensemble $\mathfrak{D}$,

$$\begin{aligned}
\rho_{\text{dist}} &= \operatorname{tr}_{\mathsf{XY}}\big(B \cdot \mathop{\mathbb{E}}_{C} \big[\boldsymbol{C}^\dagger \Pi_{\mathsf{X}}^{\text{c-dist}} \boldsymbol{C}\big] \otimes \Pi_{\mathsf{Y}}^{\text{c-dist}} \cdot |\Psi\rangle\langle\Psi|\big) && \text{(up to relative error } \delta\text{, from (1.))} \\
&= \operatorname{tr}_{\mathsf{XY}}\big(B \cdot \mathbb{1}_{\mathsf{X}} \otimes \Pi_{\mathsf{Y}}^{\text{c-dist}} \cdot |\Psi\rangle\langle\Psi|\big) + \Delta \equiv \rho_a + \Delta, && \text{(from (2.), (3.), (4.))}
\end{aligned}$$

where $\|\Delta\|_1 \leq \delta \operatorname{tr}\big(B \cdot \mathbb{1}_{\mathsf{X}} \otimes \Pi_{\mathsf{Y}}^{\text{c-dist}} \cdot |\Psi\rangle\langle\Psi|\big) \leq \delta$ from (2.), (3.), (4.). We have $\|\rho_{\text{dist}} - \rho_a\|_1 \leq 3\delta$, since the 1-norm error is upper bounded by twice the relative error [30].

This completes our proof: We have $\|\rho - \rho_a\|_1 \leq 2\delta + 3\delta$ when C is drawn from any 2-design $\mathfrak{D}$. Since the Haar ensemble is a 2-design, we have $\|\rho - \rho^H\|_1 \leq \|\rho^H - \rho_a\|_1 + \|\rho - \rho_a\|_1 \leq 4\delta + 6\delta = 10k^2/D$. □

Proof of (5.)—We proceed by induction. Let $\rho_j = \operatorname{tr}_{\mathsf{ABXY}_{>j}}(|\Psi\rangle\langle\Psi|)$ denote the reduced density matrix on on $\mathsf{XY}_{\leq j}$. The key property we use is that ρ_j is maximally mixed on Y_j.

Base case ($j = 1$): Here, $B = \mathbb{1} \otimes \mathbb{1}$ and $\Pi_{\mathsf{Y}_{\leq 1}}^{\text{c-dist}} = \mathbb{1}_{\mathsf{Y}_1}$. Hence, $\operatorname{tr}(B \cdot \Pi_{\mathsf{Y}_{\leq 1}}^{\text{c-dist}} \cdot |\Psi\rangle\langle\Psi|) = 1$.

Inductive step: We can first expand,

$$B_{\mathsf{XY}} \Pi_{\mathsf{Y}}^{\text{c-dist}} = \sum_{\substack{t,t' \in \{0,1\}^j \\ |t| = |t'|}} D^{2(j-|t|)} \cdot \left(|t', t'\rangle\langle t, t|_{\mathsf{X_C Y_C}} \otimes \sum_{\pi \in S_{|t|}} \Big[\Big(\pi_{t \to t'} \Pi_{\mathsf{Y}_{\mathsf{S},t}}^{\text{dist}} \otimes \pi_{t \to t'}\Big) \otimes |\Psi_{\text{Bell}}^{\bar{t}'}\rangle\langle\Psi_{\text{Bell}}^{\bar{t}}|\Big]_{\mathsf{X_S Y_S}} \right).$$

We now perform the trace over Y_j. Since ρ_j is maximally mixed on Y_j, the trace sets $t_j = t'_j$. If $t_j = 0$, we find

$$\operatorname{tr}_{\mathsf{XY}_j}\big(\big[\, |0,0\rangle\langle 0,0| \otimes |\Psi_{\text{Bell}}^j\rangle\langle\Psi_{\text{Bell}}^j| \,\big]_{\mathsf{XY}_j} \rho_j\big) = p_j(0) \cdot \rho_{j-1}^{(0)} / D^2, \tag{C.16}$$

where $p_j(0)\rho_{j-1}^{(0)} \equiv \mathrm{tr}_{\mathsf{XY}_j}(|0\rangle\langle 0|_{\mathsf{X_{C},}j}\,\rho_j)$ is the probability of measuring a 0 on register $\mathsf{X_{C},}j$, multiplied by the resulting normalized quantum state. Meanwhile, if $t_j = 1$, we find

$$\mathrm{tr}_{\mathsf{XY}_j}\Big(\big[\,|1,1\rangle\langle 1,1|\otimes\sum_{\pi\in S_{|t|}}\Big(\pi_{t\to t'}\otimes\pi_{t\to t'}\Pi^{\mathrm{dist}}_{\mathsf{Y_{S},}t}\Big)\,\big]_{\mathsf{XY}_j}\rho_j\,.\Big), \tag{C.17}$$

Leveraging the identity, $\mathrm{tr}_{\mathsf{Y_{S},}j}(\pi_{\mathsf{Y_S}}\Pi^{\mathrm{dist}}_{\mathsf{Y_S}}) = (D-j+1)\pi_{\mathsf{Y_{S},}<j}\Pi^{\mathrm{dist}}_{\mathsf{Y_{S},}<j}$ if $\pi(j)=j$, and $\mathrm{tr}_{\mathsf{Y_{S},}j}(\pi_{\mathsf{Y_S}}\Pi^{\mathrm{dist}}_{\mathsf{Y_S}}) = 0$ otherwise, we can compute

$$\mathrm{tr}_{\mathsf{XY}_j}\Big(\big[\,|1,1\rangle\langle 1,1|\otimes\sum_{\pi\in S_{|t|}}\Big(\pi_{t\to t'}\Pi^{\mathrm{dist}}_{\mathsf{Y_{S},}t}\otimes\pi_{t\to t'}\Big)\,\big]_{\mathsf{XY}_j}\rho_j\Big) = \frac{(D-j+1)}{D}\cdot p_j(1)\cdot\rho^{(1)}_{j-1}, \tag{C.18}$$

where $p_j(1)\rho^{(1)}_{j-1} \equiv \mathrm{tr}_{\mathsf{XY}_j}(|1\rangle\langle 1|_{\mathsf{X_{C},}j}\,\rho_j)$. We have $p_j(0)+p_j(1)=1$ by definition.

Iterating k times and summing over each $t\in\{0,1\}^k$, we find

$$\mathrm{tr}\big(B_{\mathsf{XY}}\Pi^{\text{c-dist}}_{\mathsf{Y}}\,|\Psi\rangle\langle\Psi|\big) = \sum_{t\in\{0,1\}^k}(\mathfrak{D}_{|t|}/D^{|t|})\cdot p(t) \tag{C.19}$$

where we let $p(t)$ denote the product of each $p_j(0)$ or $p_j(1)$ encountered in the iteration, and we let $\mathfrak{D}_{|t|} = D!/(D-|t|+1)!$ denote the dimension of the distinct subspace on t copies. We have $\mathrm{tr}\big(B_{\mathsf{XY}}\Pi^{\text{c-dist}}_{\mathsf{Y}}\,|\Psi\rangle\langle\Psi|\big)\le 1$, because $\sum_t p(t)=1$ and $\mathfrak{D}_{|t|}\le D^{|t|}$. We also have $\mathrm{tr}\big(B_{\mathsf{XY}}\Pi^{\text{c-dist}}_{\mathsf{Y}}\,|\Psi\rangle\langle\Psi|\big)\ge 1-k^2/D$ because $\mathfrak{D}_{|t|}\ge D^{|t|}(1-k^2/D)$. □

C.3 Proof of Theorem 11: Controlled-LRFC ensemble

We follow the same notation as in the previous section. We denote $\boldsymbol{L} = (O_L^c)^{\otimes k}$, $\boldsymbol{F} = (F^c)^{\otimes k}$, $\boldsymbol{R} = (S_R^c)^{\otimes k}$, $\boldsymbol{D} = (D^c)^{\otimes k}$. To write the expected output state in a more compact form, we let

$$B'_{\mathsf{XY}} \equiv (2D)^{2k}\cdot\mathop{\mathbb{E}}_{f_R,f,f_L}\big[(\boldsymbol{R}^\dagger\boldsymbol{F}^\dagger\boldsymbol{L}^\dagger\otimes\mathbb{1})_{\mathsf{XY}}\,|\Psi_{\mathrm{Bell}}\rangle\langle\Psi_{\mathrm{Bell}}|_{\mathsf{XY}}\,(\boldsymbol{LFR}\otimes\mathbb{1})_{\mathsf{XY}}\big] \tag{C.20}$$

denote the Bell projector twirled over O_R, F, O_L, multiplied by $(2D)^{2k}$. This yields,

$$\rho_{\mathsf{AB}} \equiv \mathop{\mathbb{E}}_{f_R,f,f_L,C}\big[\,|\psi_U\rangle\langle\psi_U|_{\mathsf{AB}}\,\big] = \mathop{\mathbb{E}}_{C}\big[\,\mathrm{tr}_{\mathsf{XY}}\left(B'_{\mathsf{XY}}\boldsymbol{C}_{\mathsf{X}}\,|\Psi\rangle\langle\Psi|_{\mathsf{ABXY}}\,\boldsymbol{C}_{\mathsf{X}}^\dagger\right)\big]. \tag{C.21}$$

We let $\rho^H_{\mathsf{AB}} \equiv \mathbb{E}_{U\sim H}\,|\psi_U\rangle\langle\psi_U|_{\mathsf{AB}}$ denote the output of the same experiment when U_c is controlled Haar-random. We will show that $\|\rho-\rho^H\|_1\le 4\delta$, where $\delta = k^2/2^n$. This establishes security against any sub-exponential time quantum adversary.

We begin by writing down several definitions and facts. We define the controlled distinct subspace projectors,

$$\Pi^{\text{c-dist}}_{\mathsf{X_L}} = \sum_{t\in\{0,1\}^k}|t\rangle\langle t|_{\mathsf{X_{L,C}}}\otimes\big[\Pi^{\mathrm{dist}}_t\otimes\mathbb{1}_{\bar t}\big]_{\mathsf{X_{L,S}}}, \qquad \Pi^{\text{c-dist}}_{\mathsf{Y_R}} = \sum_{t\in\{0,1\}^k}|t\rangle\langle t|_{\mathsf{Y_{R,C}}}\otimes\big[\Pi^{\mathrm{dist}}_t\otimes\mathbb{1}_{\bar t}\big]_{\mathsf{Y_{R,S}}}, \tag{C.22}$$

as well as their product,

$$\Pi^{\text{c-dist}}_{\mathsf{X_L Y_R}} \equiv \Pi^{\text{c-dist}}_{\mathsf{X_L}}\otimes\Pi^{\text{c-dist}}_{\mathsf{Y_R}}. \tag{C.23}$$

Here, $\mathsf{X_{L,C}}$ denotes the k control registers in $\mathsf{X_L}$, and $\mathsf{X_{L,S}}$ the k system register, and similar for $\mathsf{Y_R}$. Here, $\Pi^{\mathrm{dist}}_t = \sum_{x\in\mathrm{dist}}|x\rangle\langle x|$ is the projector onto the distinct subspace on the system registers of copies j with $t_j=1$, where sum runs over all distinct bitstrings in $[D_L]^{\otimes|t|}$ or $[D_R]^{\otimes|t|}$. Also, $\mathbb{1}_{\bar t}$ is the identity matrix on the copies j with $t_j=0$.

In analogy to B' above, we also define the operator,

$$B_{\mathsf{XY}} = D^{2(k-|t|)} \sum_{\substack{t,t' \in \{0,1\}^k \\ |t|=|t'|}} \left(\big|t',t'\big\rangle\!\big\langle t,t\big|_{\mathsf{X_C Y_C}} \otimes \sum_{\pi \in S_{|t|}} \Big[(\pi_{t\to t'} \otimes \pi_{t\to t'}) \otimes |\Psi^{\bar{t}'}_{\text{Bell}}\rangle\!\langle\Psi^{\bar{t}}_{\text{Bell}}| \Big]_{\mathsf{X_S Y_S}} \right), \quad \text{(C.24)}$$

where $\pi_{t\to t'}$ is a permutation mapping the system registers of the copies of X where $t_j = 1$ to the those where $t'_j = 1$, and similarly for Y. With these definitions in hand, our proof uses five facts:

(1) B and B' are equal on the distinct subspace, $B'_{\mathsf{XY}} \cdot \Pi^{\text{c-dist}}_{\mathsf{X_L Y_R}} = B_{\mathsf{XY}} \cdot \Pi^{\text{c-dist}}_{\mathsf{X_L Y_R}}$. To see this, we express the left side as

$$B'_{\mathsf{XY}} = \sum_{\substack{t,t' \in \{0,1\}^k \\ |t|=|t'|}} D^{2(k-|t|)} \cdot \big|t,t\big\rangle\!\big\langle t',t'\big|_{\mathsf{X_C Y_C}} \otimes \Big[(\tilde{\pi}_{t\to t'} \otimes \tilde{\pi}_{t\to t'}) \cdot \Big(B'_t \otimes |\Psi^{\bar{t}}_{\text{Bell}}\rangle\!\langle\Psi^{\bar{t}}_{\text{Bell}}| \Big) \Big]_{\mathsf{X_S Y_S}}, \quad \text{(C.25)}$$

where $\tilde{\pi}_{t\to t'}$ is any permutation that maps copies j where $t_j = 1$ to those where $t'_j = 1$, and we abbreviate

$$(B'_t)_{\mathsf{XY}_t} \equiv D^{2|t|} \mathop{\mathbb{E}}_{f_R,f,f_L} \big[(O_R^{\dagger,\otimes|t|} F^{\dagger,\otimes|t|} O_L^{\dagger,\otimes|t|} \otimes \mathbb{1})_{\mathsf{XY}_t} \big|\Psi^t_{\text{Bell}}\big\rangle\!\big\langle\Psi^t_{\text{Bell}}\big|_{\mathsf{XY}_t} (O_L^{\otimes|t|} F^{\otimes|t|} O_R^{\otimes|t|} \otimes \mathbb{1})_{\mathsf{XY}_t} \big] \quad \text{(C.26)}$$

We can perform precisely the same decomposition for B, replacing $(B'_t)_{\mathsf{XY}_t}$ with $(B_t)_{\mathsf{XY}_t} \equiv \sum_{\pi\in S_{|t|}} \pi \otimes \pi$. Inserting the distinct subspace projectors, we find

$$\begin{aligned} B'_t(\Pi^{\text{dist}}_{\mathsf{X}_{\mathsf{L}},t} \otimes \Pi^{\text{dist}}_{\mathsf{Y}_{\mathsf{R}},t}) &= D^{|t|} \mathop{\mathbb{E}}_{f_R,f,f_L} \sum_{\substack{x_L,y_R \\ \text{dist}}} \sum_{\substack{\tilde{x}_L,\tilde{y}_R \\ \text{dist}}} (-1)^{f(x_L\|y_R)+f(\tilde{x}_L\|\tilde{y}_R)} \, |x_L\|y_R \oplus f_R(x_L)\rangle_{\mathsf{X}_t} \otimes \\ &\qquad |x_L \oplus f_L(y_R)\|y_R\rangle_{\mathsf{Y}_t} \langle \tilde{x}_L\|\tilde{y}_R \oplus f_R(\tilde{x}_L)|_{\mathsf{X}_t} \otimes \langle \tilde{x}_L \oplus f_L(\tilde{y}_R)\|\tilde{y}_R|_{\mathsf{Y}_t} \\ &= D^{2|t|} \mathop{\mathbb{E}}_{f_R,f_L} \sum_{\substack{x_L,y_R \\ \text{dist}}} \sum_{\pi\in S_{|t|}} (|x_L\|y_R \oplus f_R(x_L)\rangle\!\langle x_L\|y_R \oplus f_R(x_L)|\,\pi)_{\mathsf{X}_t} \otimes \\ &\qquad (|x_L \oplus f_L(y_R)\|y_R\rangle\!\langle x_L \oplus f_L(y_R)\|y_R|\,\pi)_{\mathsf{Y}_t}, \end{aligned}$$

where in the second line, taking the expectation over f enforces $x_L\|y_R = \pi(x_L\|y_R)$ for some $\pi \in S_{|t|}$. Taking the expectation over f_L and f_R yields,

$$B'_t(\Pi^{\text{dist}}_{\mathsf{X}_{\mathsf{L}},t} \otimes \Pi^{\text{dist}}_{\mathsf{Y}_{\mathsf{R}},t}) = \sum_{\substack{x_L,y_R \\ \text{dist}}} \sum_{y_L,x_R} (|x_L\|x_R\rangle\!\langle x_L\|x_R|\,\pi)_{\mathsf{X}_t} \otimes (|y_L\|y_R\rangle\!\langle y_L\|y_R|\,\pi)_{\mathsf{Y}_t}, \quad \text{(C.27)}$$

since $y_R \oplus f_R(x_L)$ is a uniformly random bitstring in $[D_R]^{\otimes|t|}$ when x_L is distinct, and similar for $x_L \oplus f_L(y_R)$. This is clearly equal to $B_t(\Pi^{\text{dist}}_{\mathsf{X}_{\mathsf{L}},t} \otimes \Pi^{\text{dist}}_{\mathsf{Y}_{\mathsf{R}},t})$, using the resolution of the identity, $\mathbb{1}_{\mathsf{XY}_t} = \sum_{x_L,x_R,y_L,y_R} |x_L\|x_R\rangle\!\langle x_L\|x_R|_{\mathsf{X}_t} \otimes |y_L\|y_R\rangle\!\langle y_L\|y_R|_{\mathsf{Y}_t}$.

(2.) The twirl of $\Pi^{\text{c-dist}}_{\mathsf{X_L}}$ over a controlled 2-design is close to the identity, $\big\|\mathbb{1} - \mathbb{E}_{C\sim\mathfrak{D}} \big[\boldsymbol{C}_{\mathsf{X}}{}^\dagger \Pi^{\text{c-dist}}_{\mathsf{X_L}} \boldsymbol{C}_{\mathsf{X}}\big]\big\|_\infty \leq \delta_L$. To show this, we bound $\mathbb{1} - \Pi^{\text{c-dist}}_{\mathsf{X_L}} \leq \sum_{t\in\{0,1\}^k} \sum_{i<j\in t} |t\rangle\!\langle t|_{\mathsf{X}_{\mathsf{L,C}}} \otimes \Pi^{ij}_{\mathsf{X}_{\mathsf{L,S}}}$, where

$$\Pi^{\text{eq},ij}_{\mathsf{X}_{\mathsf{L,S}}} = \sum_{x\in[D_L]} |x,x\rangle\!\langle x,x|_{\mathsf{X}_{\mathsf{L_S},i}\mathsf{X}_{\mathsf{L_S},j}}. \quad \text{(C.28)}$$

The twirl over a controlled 2-design yields $\mathbb{E}_C[\boldsymbol{C}_{\mathsf{X}}{}^\dagger \sum_t \sum_{i<j\in t} |t\rangle\!\langle t| \otimes \Pi^{\text{eq},ij}_{\mathsf{X}_{\mathsf{L,S}}} \boldsymbol{C}_{\mathsf{X}}] = \sum_t \sum_{i<j\in t} |t\rangle\!\langle t| \otimes (a\mathbb{1} + b\mathcal{S}_{ij})$, where $\mathcal{S}_{ij}$ is the swap operator, and $a = (D_R^2 D_L - 1)/(D_R^2 D_L^2 - 1) \leq 1/D_L$ and $b =$

$(D_R D_L - D_R)/(D_R^2 D_L^2 - 1) \leq 1/D_R D_L$. Hence, the spectral norm is less $\max_t \sum_{i<j\in t}(a+b) \leq k^2(a+b)/2 \leq \delta_L(1+1/D_R) \leq \delta_L$.

(3.) B and B' commute with $\Pi^{\text{c-dist}}_{\mathsf{X_L}}$ and $\Pi^{\text{c-dist}}_{\mathsf{Y_R}}$. B also commutes with $\boldsymbol{C}$.

(4.) For any positive Hermitian matrices P, Q, R, where P and Q commute, we have $\operatorname{tr}(PQR) \leq \|Q\|_\infty \operatorname{tr}(PR)$. This follows from Holder's inequality,

$$\operatorname{tr}(PQR) = \operatorname{tr}(Q\sqrt{P}R\sqrt{P}) \leq \|Q\|_\infty \|\sqrt{P}R\sqrt{P}\|_1 = \|Q\|_\infty \operatorname{tr}(PR).$$

(5.) We have $1 - \delta_R \leq \operatorname{tr}\Big(B_{\mathsf{XY}} \Pi^{\text{c-dist}}_{\mathsf{Y_R}} |\Psi\rangle\langle\Psi|\Big) \leq 1$. This is shown below the current proof.

We can now prove the claim. We insert $\mathbb{1} = \Pi^{\text{c-dist}}_{\mathsf{X_L Y_R}} + (\mathbb{1} - \Pi^{\text{c-dist}}_{\mathsf{X_L Y_R}})$ to decompose ρ into two terms,

$$\rho = \rho_{\text{dist}} + \delta\rho \equiv \mathop{\mathbb{E}}_{C} \big[\operatorname{tr}_{\mathsf{XY}}(B' \cdot \Pi^{\text{c-dist}}_{\mathsf{X_L Y_R}} \cdot \boldsymbol{C}\,|\Psi\rangle\langle\Psi|\,\boldsymbol{C}^\dagger)\big] + \mathop{\mathbb{E}}_{C} \big[\operatorname{tr}_{\mathsf{XY}}(B' \cdot (\mathbb{1} - \Pi^{\text{c-dist}}_{\mathsf{X_L Y_R}}) \cdot \boldsymbol{C}\,|\Psi\rangle\langle\Psi|\,\boldsymbol{C}^\dagger)\big].$$

We bound the second term as follows,

$$\begin{aligned}
\|\delta\rho\|_1 = \operatorname{tr}(\delta\rho) &= \mathop{\mathbb{E}}_{C} \Big[\operatorname{tr}\Big(B'(\mathbb{1} - \Pi^{\text{c-dist}}_{\mathsf{X_L Y_R}})\boldsymbol{C}\,|\Psi\rangle\langle\Psi|\,\boldsymbol{C}^\dagger\Big)\Big] && \text{(since } \delta\rho \text{ is positive)} \\
&= 1 - \mathop{\mathbb{E}}_{C} \big[\operatorname{tr}\big(B \cdot \Pi^{\text{c-dist}}_{\mathsf{X_L Y_R}} \cdot \boldsymbol{C}\,|\Psi\rangle\langle\Psi|\,\boldsymbol{C}^\dagger\big)\big], && \text{(from } \operatorname{tr}(\rho) = 1 \text{ and (1.))} \\
&= 1 - \operatorname{tr}\big(B \cdot \mathop{\mathbb{E}}_{C} \big[\boldsymbol{C}^\dagger \Pi^{\text{c-dist}}_{\mathsf{X_L}} \boldsymbol{C}\big] \otimes \Pi^{\text{c-dist}}_{\mathsf{Y_R}} \cdot |\Psi\rangle\langle\Psi|\big), && \text{(from (3.))} \\
&\leq 1 - (1-\delta_L) \operatorname{tr}\big(B(\mathbb{1} \otimes \Pi^{\text{c-dist}}_{\mathsf{Y_R}})\,|\Psi\rangle\langle\Psi|\big) && \text{(from (2.), (3.), (4.))} \\
&= 1 - (1-\delta_L)(1-\delta_R) \leq \delta_L + \delta_R. && \text{(from (5.))}
\end{aligned}$$

Meanwhile, the first term is close to a fixed density matrix ρ_a, independent of the ensemble $\mathfrak{D}$,

$$\begin{aligned}
\rho_{\text{dist}} &= \operatorname{tr}_{\mathsf{XY}}\big(B \cdot \mathop{\mathbb{E}}_{C} \big[\boldsymbol{C}^\dagger \Pi^{\text{c-dist}}_{\mathsf{X_L}} \boldsymbol{C}\big] \otimes \Pi^{\text{c-dist}}_{\mathsf{Y_R}} \cdot |\Psi\rangle\langle\Psi|\big) && \text{(from (1.))} \\
&= \operatorname{tr}_{\mathsf{XY}}\big(B \cdot \mathbb{1}_{\mathsf{X_L}} \otimes \Pi^{\text{c-dist}}_{\mathsf{Y_R}} \cdot |\Psi\rangle\langle\Psi|\big) + \Delta \equiv \rho_a + \Delta, && \text{(from (2.), (3.), (4.))}
\end{aligned}$$

where $\|\Delta\|_1 \leq \delta_L \operatorname{tr}\Big(B \cdot \mathbb{1}_{\mathsf{X_L}} \otimes \Pi^{\text{c-dist}}_{\mathsf{Y_R}} \cdot |\Psi\rangle\langle\Psi|\Big) \leq \delta_L$ from (2.), (3.), (4.).

This completes our proof: We have $\|\rho - \rho_a\|_1 \leq 2\delta_L + \delta_R$ when C is drawn from any 2-design $\mathfrak{D}$. Since the Haar ensemble is a 2-design, we have $\|\rho - \rho^H\|_1 \leq \|\rho^H - \rho_a\|_1 + \|\rho - \rho_a\|_1 \leq 4\delta_L + 2\delta_R = 4k^2/D_L + 2k^2/D_R$. □

Proof of (5.)—We proceed by induction, similar to the previous section. Let $\rho_j = \operatorname{tr}_{\mathsf{ABXY}_{>j}}(|\Psi\rangle\langle\Psi|)$ denote the reduced density matrix on on $\mathsf{XY}_{\mathsf{R},\leq j}$. The key property we use is that ρ_j is maximally mixed on Y_j.

Base case ($j = 1$): Here, $B = \mathbb{1} \otimes \mathbb{1}$ and $\Pi^{\text{c-dist}}_{\mathsf{Y_{R,\leq 1}}} = \mathbb{1}_{\mathsf{Y_{R,1}}}$. Hence, $\operatorname{tr}\Big(B \cdot \Pi^{\text{c-dist}}_{\mathsf{Y_{R,\leq 1}}} \cdot |\Psi\rangle\langle\Psi|\Big) = 1$.

Inductive step: We can first expand,

$$B_{\mathsf{XY}} \Pi^{\text{c-dist}}_{\mathsf{Y_R}} = \sum_{\substack{t,t'\in\{0,1\}^j \\ |t|=|t'|}} D^{2(j-|t|)} \cdot \left(|t', t'\rangle\langle t, t|_{\mathsf{X_C Y_C}} \otimes \sum_{\pi\in S_{|t|}} \Big[\Big(\pi_{t\to t'} \Pi^{\text{dist}}_{\mathsf{Y_{R,S,t}}} \otimes \pi_{t\to t'}\Big) \otimes |\Psi^{\bar{t}'}_{\text{Bell}}\rangle\langle\Psi^{\bar{t}}_{\text{Bell}}|\Big]_{\mathsf{X_S Y_S}} \right).$$

We now perform the trace over Y_j. Since ρ_j is maximally mixed on Y_j, the trace sets $t_j = t'_j$. If $t_j = 0$, we find

$$\mathrm{tr}_{\mathsf{X}\mathsf{Y}_j}\big(\big[\,|0,0\rangle\langle 0,0| \otimes |\Psi^j_{\mathrm{Bell}}\rangle\langle\Psi^j_{\mathrm{Bell}}|\,\big]_{\mathsf{X}\mathsf{Y}_j}\rho_j\big) = p_j(0)\cdot\rho^{(0)}_{j-1}/D^2, \tag{C.29}$$

where $p_j(0)\rho^{(0)}_{j-1} \equiv \mathrm{tr}_{\mathsf{X}\mathsf{Y}_j}(|0\rangle\langle 0|_{\mathsf{X}_{\mathsf{C},j}}\,\rho_j)$ is the probability of measuring a 0 on register $\mathsf{X}_{\mathsf{C},j}$, multiplied by the resulting normalized quantum state. Meanwhile, if $t_j = 1$, we find

$$\mathrm{tr}_{\mathsf{X}\mathsf{Y}_j}\Big(\big[\,|1,1\rangle\langle 1,1| \otimes \sum_{\pi\in S_{|t|}}\Big(\pi_{t\to t'}\Pi^{\mathrm{dist}}_{\mathsf{Y}_{\mathsf{R},\mathsf{S},\mathsf{t}}} \otimes \pi_{t\to t'}\Big)\,\big]_{\mathsf{X}\mathsf{Y}_j}\rho_j\,.\Big), \tag{C.30}$$

Leveraging the identity,

$$\mathrm{tr}_{\mathsf{Y}_{\mathsf{S},j}}(\pi_{\mathsf{Y}_\mathsf{S}}\Pi^{\mathrm{dist}}_{\mathsf{Y}_{\mathsf{R},\mathsf{S}}}) = D_L(D_R - j + 1)\pi_{\mathsf{Y}_{\mathsf{S},<j}}\Pi^{\mathrm{dist}}_{\mathsf{Y}_{\mathsf{R},\mathsf{S},<j}}$$

if $\pi(j) = j$, and $\mathrm{tr}_{\mathsf{Y}_{\mathsf{S},j}}(\pi_{\mathsf{Y}_\mathsf{S}}\Pi^{\mathrm{dist}}_{\mathsf{Y}_{\mathsf{R},\mathsf{S}}}) = 0$ otherwise, we can compute

$$\mathrm{tr}_{\mathsf{X}\mathsf{Y}_j}\Big(\big[\,|1,1\rangle\langle 1,1| \otimes \sum_{\pi\in S_{|t|}}\Big(\pi_{t\to t'}\Pi^{\mathrm{dist}}_{\mathsf{Y}_{\mathsf{R},\mathsf{S},\mathsf{t}}} \otimes \pi_{t\to t'}\Big)\,\big]_{\mathsf{X}\mathsf{Y}_j}\rho_j\Big) = \frac{D_L(D_R - j + 1)}{D_L D_R}\cdot p_j(1)\cdot\rho^{(1)}_{j-1}, \tag{C.31}$$

where $p_j(1)\rho^{(1)}_{j-1} \equiv \mathrm{tr}_{\mathsf{X}\mathsf{Y}_j}(|1\rangle\langle 1|_{\mathsf{X}_{\mathsf{C},j}}\,\rho_j)$. We have $p_j(0) + p_j(1) = 1$ by definition.

Iterating k times and repeating for every $t \in \{0,1\}^k$, we find

$$\mathrm{tr}\big(B_{\mathsf{X}\mathsf{Y}}\Pi^{\text{c-dist}}_{\mathsf{Y}_\mathsf{R}}\,|\Psi\rangle\langle\Psi|\big) = \sum_{t\in\{0,1\}^k}(\mathfrak{D}^{|t|}_R/D^{|t|}_R)\cdot p(t) \tag{C.32}$$

where we let $p(t)$ denote the product of each $p_j(0)$ or $p_j(1)$ encountered in the iteration, and we let $\mathfrak{D}^{|t|}_R = ((D_R)!/(D_R - |t| + 1)!)$ denote the dimension of the distinct subspace on $[D_R]^{\otimes|t|}$. We have $\mathrm{tr}\Big(B_{\mathsf{X}\mathsf{Y}}\Pi^{\text{c-dist}}_{\mathsf{Y}_\mathsf{R}}\,|\Psi\rangle\langle\Psi|\Big) \le 1$, because $\sum_t p(t) = 1$ and $\mathfrak{D}^{|t|}_R \le D^{|t|}_R$. We also have $\mathrm{tr}\Big(B_{\mathsf{X}\mathsf{Y}}\Pi^{\text{c-dist}}_{\mathsf{Y}_\mathsf{R}}\,|\Psi\rangle\langle\Psi|\Big) \ge 1 - k^2/D_R$ because $\mathfrak{D}^{|t|}_R \ge D^{|t|}_R(1 - k^2/D_R)$. □

D Translation-invariant random unitaries

In this section, we present the extension of our results to translation-invariant random unitaries. For simplicity, we focus on random unitaries that are invariant with respect to translations over the same distance, 2ξ, as the patch size in our two-layer brickwork construction. We begin in Appendix D.1 by establishing a few basic results on translation-invariant random unitaries. We then analyze the gluing construction for translation-invariant unitaries in Appendix D.2. We show how this leads to extremely low depth translation-invariant PRUs. In Appendix D.3, we extend our results to translation-invariant unitaries that also have on-site symmetries.

D.1 Approximate translation-invariant Haar twirl

We consider a periodic 1D chain of n qubits. We suppose that the system is symmetric under translations over a distance 2ξ. We let $\mathcal{T}$ denote the operator which implements this translation. We have $\mathcal{T}^m = 1$, where $m = n/2\xi$.

We are interested in the Haar measure over unitaries that commute with the translation symmetry operator. To spare time, we will not analyze the exact expression involving Weingarten functions for

the translation-invariant Haar twirl. Instead, we will begin by positing the following approximate expression for the translation-invariant Haar twirl,

$$\Phi_a^{\mathrm{TI}}(\rho) = \frac{1}{D^k} \sum_{\pi} \sum_{\boldsymbol{t}} \mathrm{tr}\left(\rho \cdot \pi^{-1} \cdot \bigotimes_{j=1}^{k} \mathcal{T}^{-t_j}\right) \cdot \bigotimes_{j=1}^{k} \mathcal{T}^{t_j} \cdot \pi. \tag{D.1}$$

Here, $\pi \in S_k$ is summed over all permutations amongst the k copies, and $\boldsymbol{t} = (t_1, \ldots, t_k)$ is summed over all translations, over distance $2\xi t_j$, in each copy j.

The key technical result of this subsection is the following lemma. The lemma allows one to bound the relative error to the *exact* translation-invariant Haar twirl, by bounding the additive error to the *approximate* translation-invariant Haar twirl on the EPR state.

Lemma 5 (Translation-invariant unitary designs from EPR states)**.** *An ensemble $\mathcal{E}$ of translation-invariant unitaries forms a translation-invariant ε-approximate unitary k-design with relative error*

$$\varepsilon = \frac{3D^{2k}}{n^k k!} \left\| [\delta\Phi \otimes \mathbb{1}](P_{EPR}) \right\|_\infty, \tag{D.2}$$

where $\delta\Phi = \Phi_{\mathcal{E}} - \Phi_a^{\mathrm{TI}}$, and P_{EPR} is the EPR state on $\mathcal{H}^{\otimes k} \otimes \mathcal{H}^{\otimes k}$.

Curiously, the lemma can be proven even without knowing that Φ_a^{TI} is close to the translation-invariant Haar twirl in relative error. The only property of Φ_a^{TI} that we require is that it is invariant under the Haar twirl, $\Phi_H^{\mathrm{TI}} \circ \Phi_a^{\mathrm{TI}} = \Phi_a^{\mathrm{TI}}$. This is convenient, as it allows us to sidestep a brute force Weingarten analysis of the exact twirl.

Proof. The first part of our proof proceeds exactly analogously to the proof of Lemma 2 in Ref. [30] and the proof of Lemma 3 in this work. The sole difference is that the non-zero eigenvalue of $[\Phi_a^{\mathrm{TI}} \otimes \mathbb{1}](P_{\mathrm{EPR}})$ is now equal to $n^k k!/D^{2k}$, owing to the $n^k k!$ permutations appearing Φ_a^{TI}. This immediately shows that $\Phi_{\mathcal{E}}$ is close to the *approximate* Haar twirl,

$$(1-\varepsilon')\Phi_a^{\mathrm{TI}} \preceq \Phi_{\mathcal{E}} \preceq (1+\varepsilon')\Phi_a^{\mathrm{TI}}, \tag{D.3}$$

up to relative error $\varepsilon' = (D^{2k}/n^k k!) \left\| [\delta\Phi \otimes \mathbb{1}](P_{\mathrm{EPR}}) \right\|_\infty$.

To bound the relative error between $\Phi_{\mathcal{E}}$ and the *exact* Haar twirl Φ_H^{TI}, we develop the following short-cut. We first apply the exact Haar twirl to each term in Eq. (D.3). Using $\Phi_H^{\mathrm{TI}} \circ \Phi_{\mathcal{E}} = \Phi_H^{\mathrm{TI}}$ (from the definition of the Haar measure) and $\Phi_H^{\mathrm{TI}} \circ \Phi_a^{\mathrm{TI}} = \Phi_a^{\mathrm{TI}}$ (since $\Phi_a^{\mathrm{TI}}(\cdot)$ is a sum of operators that commute with any translation-invariant unitary), we find,

$$(1-\varepsilon')\Phi_a^{\mathrm{TI}} \preceq \Phi_H^{\mathrm{TI}} \preceq (1+\varepsilon')\Phi_a^{\mathrm{TI}}. \tag{D.4}$$

Combining this with Eq. (D.3), we have

$$(1-3\varepsilon')\Phi_H^{\mathrm{TI}} \preceq \frac{1-\varepsilon'}{1+\varepsilon'}\Phi_a^{\mathrm{TI}} \preceq \Phi_{\mathcal{E}} \preceq \frac{1+\varepsilon'}{1-\varepsilon'}\Phi_H^{\mathrm{TI}} \preceq (1+3\varepsilon')\Phi_H^{\mathrm{TI}}, \tag{D.5}$$

where the outside inequalities hold for $3\varepsilon' \leq 1$. Setting $\varepsilon = 3\varepsilon'$ completes the proof. □

D.2 Translation-invariant PRUs in extremely low depth

We now show that the two-layer brickwork construction naturally allows one to construct translation-invariant random unitaries. As before, we break the system into $2m = n/\xi$ patches of ξ qubits each, and apply a translation-invariant version of the two-layer circuit,

$$U = U_2^{\otimes m} \cdot U_1^{\otimes m}, \tag{D.6}$$

where the first layer is a tensor product of m applications of the unitary U_1 acting on all even nearest-neighbor pairs of patches, and the second layer is a tensor product of m applications of U_2 acting on all odd nearest-neighbor pairs of patches. The unitary is clearly symmetric under translations, $\mathcal{T}U = U\mathcal{T}$.

Our main result is the following theorem, which shows that U forms a translation-invariant random unitary design.

Theorem 12 (Gluing translation-invariant random unitaries)**.** *For any $\varepsilon \leq 1$. Consider the translation-invariant two-layer brickwork circuit, U [Eq. (D.6)]. Suppose that the small random unitaries, U_1 and U_2, are each drawn from $\frac{\varepsilon}{4}$-approximate unitary mk-designs. Then U forms an ε-approximate translation-invariant unitary k-design, as long as $\xi \geq \log_2(32n^6k^6/\varepsilon)$.*

By drawing each of U_1 and U_2 from a PRU ensemble on 2ξ qubits with security against sub-exponential time quantum adversaries, the theorem immediately leads to translation-invariant pseudorandom unitaries in extremely low depth.

Corollary 2 (Translation-invariant PRUs in extremely low depth)**.** *Translation-invariant PRUs exist, and can be formed in depth* $\mathrm{poly}(\log n)$ *in geometrically-local circuits, and* $\mathrm{poly}\log\log n$ *in all-to-all connected circuits.*

Proof of Theorem 12. We let $K \equiv mk$, and set $D \equiv 2^\xi$ for the duration of this proof. By assumption, the twirl when U_1 and U_2 are drawn from approximate designs is close to the twirl when they are both drawn from the Haar measure, up to relative error $(1+\varepsilon/4)(1+\varepsilon/4)$. The latter is given by,

$$\begin{aligned}[\Phi_{\mathcal{E}} \otimes \mathbb{1}](P_{\mathrm{EPR}}) &= \frac{1}{D^{3k}} \sum_{\sigma,\tilde{\sigma}\in S_K} \mathrm{tr}_e((\mathcal{T}\tilde{\sigma}\mathcal{T}^{-1})^{-1}\sigma) \cdot \mathrm{tr}_o(\tilde{\sigma}^{-1}\sigma) \cdot [(\mathcal{T}\tilde{\sigma}\mathcal{T}^{-1})_e \otimes \tilde{\sigma}_o] \otimes [\sigma_e \otimes \sigma_o] \\ &= \frac{1}{D^{2k}} \sum_{\sigma,\tilde{\sigma}\in S_K} D^{-|\tilde{\sigma}^{-1}\mathcal{T}^{-1}\sigma\mathcal{T}|} \cdot D^{-|\tilde{\sigma}^{-1}\sigma|} \cdot [(\mathcal{T}\tilde{\sigma}\mathcal{T}^{-1})_e \otimes \tilde{\sigma}_o] \otimes [\sigma_e \otimes \sigma_o],\end{aligned} \tag{D.7}$$

up to an additional relative error $(1+K^2/D^2)^2$ from replacing each small Haar twirl with its approximate form (from Lemma 2 of Ref. [30]). Here, $\sigma, \tilde{\sigma} \in S_K$ are summed over all possible permutations of the $K = mk$ patches within the k copies. We can also consider the analogous expression for the approximate Haar twirl,

$$[\Phi_a^{\mathrm{TI}} \otimes \mathbb{1}](P_{\mathrm{EPR}}) = \frac{1}{D^{2k}} \sum_{\sigma\,:\,[\sigma,\mathcal{T}^{\otimes k}]=0} [\sigma_e \otimes \sigma_o] \otimes [\sigma_e \otimes \sigma_o], \tag{D.8}$$

where σ is summed over the $n^k k!$ permutations in S_K that commute with $\mathcal{T}^{\otimes k}$. Note that this summation is precisely equivalent to the summation over π and $\boldsymbol{t}$ in Eq. (D.1), upon identifying $\sigma = \bigotimes_{j=1}^k \mathcal{T}^{t_j} \cdot \pi$. The terms in Eq. (D.8) exactly correspond to the terms in Eq. (D.7) with $\sigma = \tilde{\sigma}$ and $[\sigma, \mathcal{T}^{\otimes k}] = 0$. Taking the difference between the two expressions and applying the triangle inequality, we have,

$$\|[\delta\Phi \otimes \mathbb{1}](P_{\mathrm{EPR}})\|_\infty \leq \frac{1}{D^{2k}} \left(\sum_{\sigma,\tilde{\sigma}\in S_K} D^{-|\tilde{\sigma}^{-1}\mathcal{T}^{-1}\sigma\mathcal{T}|} D^{-|\tilde{\sigma}^{-1}\sigma|} - n^k k! \right), \tag{D.9}$$

where the negative term accounts for the $n^k k!$ permutations with $\sigma = \tilde{\sigma}$ and $[\sigma, \mathcal{T}^{\otimes k}] = 0$.

We bound the remaining sum in two steps. First, to compute the sum over $\tilde{\sigma}$, we prove the following proposition.

Proposition 4. $\sum_{\Delta\in S_K} D^{-|\Delta|}D^{-|\Delta^{-1}\tau|} \leq 8(K^2/2D)^{-|\tau|}$ *for any permutation* $\tau \in S_K$.

Proof. Let us write $\sum_{\Delta\in S_K} D^{-|\Delta|}D^{-|\Delta^{-1}\tau|} = \sum_{r_1,r_2} N(r_1,r_2)D^{-r_1-r_2}$, where $N(r_1,r_2)$ counts the number of Δ with both $|\Delta| = r_1$ and $|\Delta^{-1}\tau| = r_2$. We can upper bound,

$$N(r_1,r_2) \leq \min(N_1(r_1), N_2(r_2)),$$

where $N_1(r_1)$ counts the number of Δ with $|\Delta| = r_1$, and $N_2(r_2)$ counts the number with $|\Delta^{-1}\tau| = r_2$. We have $N_1(r_1) \leq (K^2/2)^{r_1}$, since there are at most $\binom{K}{2} \leq K^2/2$ ways to place each of r_1 transpositions, to obtain Δ from the identity permutation. Similarly, $N_2(r_2) \leq (K^2/2)^{r_2}$. Thus, we have $\sum_{\Delta\in S_K} D^{-|\Delta|}D^{-|\Delta^{-1}\tau|} \leq \sum_{r_1,r_2}(K^2/2)^{\min(r_1,r_2)}D^{-r_1-r_2}$.

To proceed, let us re-arrange the sums via $\sum_{r_1,r_2} = \sum_r \sum_{r_1+r_2=r}$. Note that for a fixed $r \equiv r_1+r_2$, there are $r+1$ valid values of r_1, r_2. We can quickly upper bound the sum over these r_1, r_2,

$$\sum_{r_1+r_2=r} (K^2/2)^{\min(r_1,r_2)} \leq 2\sum_{r_1=0}^{\lfloor r/2\rfloor} (K^2/2)^{r_1} = 2\frac{(K^2/2)^{\lfloor r/2\rfloor} - 2/K^2}{1-2/K^2} \leq 4(K^2/2)^{r/2}, \tag{D.10}$$

where the final inequality assumes $K^2 \geq 4$. This yields, $\sum_{\Delta\in S_K} D^{-|\Delta|}D^{-|\Delta^{-1}\tau|} \leq 4\sum_r (K^2/2)^{r/2}D^{-r}$. Now, we have $|\Delta^{-1}\tau| \geq |\tau| - |\Delta|$, since each transposition in Δ can decrease the distance of $\Delta^{-1}\tau$ by at most one. Thus, the sum is restricted to $r = r_1 + r_2 \geq |\tau|$. We can compute,

$$4\sum_{r=|\tau|}^{2K} (K^2/2)^{r/2}D^{-r} \leq 4\left(\frac{K^2}{2D}\right)^{|\tau|} \frac{1}{1-K^2/2D} \leq 8\left(\frac{K^2}{2D}\right)^{|\tau|}, \tag{D.11}$$

where the final inequality assumes $K^2/D \leq 1$. This completes the proof. □

Applying Proposition 4 to our spectral norm bound, with $\tau = \sigma^{-1}\mathcal{T}^{-1}\sigma\mathcal{T}$, yields,

$$\|[\delta\Phi \otimes \mathbb{1}](P_{\text{EPR}})\|_\infty \leq \frac{8}{D^{2k}} \sum_{\sigma\,:\,[\sigma,\mathcal{T}^{\otimes k}]\neq 0} (K^2/2D)^{-|\sigma^{-1}\mathcal{T}^{-1}\sigma\mathcal{T}|}. \tag{D.12}$$

As our second step, we compute the remaining sum using the following counting argument.

Proposition 5. *The number of permutations* $\sigma \in S_K$ *with* $|\sigma^{-1}\mathcal{T}^{-1}\sigma\mathcal{T}| = r$ *is upper bounded by* $(n^k k!)K^{4r}$.

Proof. Any permutation of distance r has at most $2r$ points that are not fixed. In the present context, a point $(i,j) \in [n] \otimes [k]$ is fixed if $\sigma((i,j)+(1,0)) = \sigma((i,j)) + (1,0)$. There are $\binom{K}{2r} \leq K^{2r}$ ways to choose the location of the non-fixed points (i.e. the (i,j) for which this condition is not satisfied), and at most $K!/(K-2r)! \leq K^{2r}$ locations they can be sent to by σ. Thus, there are at most K^{4r} configurations.

For the remaining, fixed, points, we note that, as soon as $\sigma((i,j))$ is decided for all the non-fixed points in a given copy j, the values of $\sigma((i,j))$ for the remaining fixed points in that copy are fully determined. This follows because the value of $\sigma((i,j))$ for each fixed point is determined by the value at $(i+1,j)$, which is determined by $(i+2,j)$, and so on, until one reaches a non-fixed point. The only exception to this argument is if all points in a given copy j are fixed. In this case, we are free to choose $\sigma((i,j))$ for a *single* point in copy j, and the remaining $n-1$ points are fully determined by this choice. Our choice must place $\sigma((i,j))$ into a new copy where no points have been placed thus far, because the fixed point condition will place all n qubits from copy j onto this new copy. There

are at most k copies with no non-fixed points. Thus, there are at most $n^k k!$ ways to place the fixed points.

Combining our two upper bounds, we find that there are at most $(n^k k!)K^{4r}$ ways to place the fixed and non-fixed points of σ. This completes our proof. □

Applying Proposition 5 to our spectral norm bound yields,

$$\|[\delta\Phi \otimes \mathbb{1}](P_{\mathrm{EPR}})\|_\infty \leq \frac{n^k k!}{D^{2k}} \cdot 8 \sum_{r=1}^{\infty} (K^6/2D)^{-r} = \frac{n^k k!}{D^{2k}} \cdot 8K^6/D, \tag{D.13}$$

where the final inequality holds for $K^6 \leq D$. Applying Lemma 5 shows that our ensemble is a translation-invariant unitary k-design with relative error $(1+\varepsilon/4)^2(1+K^2/D^2)^2(1+8K^6/D)$. We set $8K^6/D \leq \varepsilon/4$, i.e. $\xi \geq \log_2(32n^6k^6/\varepsilon)$, which also yields $K^2/D^2 \leq \varepsilon/(32DK^4) \leq \varepsilon/1088$, assuming trivially that $D, K \geq 2$. The inequality $(1+\varepsilon/4)^3(1+\varepsilon/1088)^2 \leq 1+\varepsilon$ for $\varepsilon \leq 1$ completes our proof. □

D.3 Symmetric translation-invariant PRUs in extremely low depth

We conclude our discussion of translation-invariant random unitaries by showing that the results of the previous two subsections also extend to translation-invariant random unitaries with any discrete on-site symmetry. We will move relatively quickly through the details in this section, since the derivations exactly mimic those previous.

We consider a periodic 1D chain of n qubits. We suppose that the system possesses a translation symmetry $\mathcal{T}$ over a distance 2ξ, as well as a discrete on-site symmetry $R_g = \bigotimes_{i=1}^n R_g^i$, where $g \in G$ is the symmetry group element. As in Appendix D.2, we will begin by positing an expression for the symmetric translation-invariant Haar twirl,

$$\Phi_a^{G,\mathrm{TI}}(\rho) = \frac{1}{D^k} \sum_{\pi} \sum_{\boldsymbol{t}} \sum_{\boldsymbol{g}} \mathrm{tr}\left(\rho \cdot \pi^{-1} \cdot \bigotimes_{j=1}^k \mathcal{T}^{-t_j} \cdot R_{\boldsymbol{g}}^{-1}\right) \cdot R_{\boldsymbol{g}} \cdot \bigotimes_{j=1}^k \mathcal{T}^{t_j} \cdot \pi. \tag{D.14}$$

The expression is precisely analogous to our previous formulas for the approximate symmetric and translation-invariant Haar twirls. Next, we can establish an analogous technical lemma for bounding the relative error of any symmetric translation-invariant unitary design.

Lemma 6 (Symmetric translation-invariant unitary designs from EPR states)**.** *For any discrete on-site symmetry group G. An ensemble $\mathcal{E}$ of symmetric translation-invariant unitaries forms a symmetric translation-invariant ε-approximate unitary k-design with relative error*

$$\varepsilon = \frac{3D^{2k}}{|G|^k n^k k!} \|[\delta\Phi \otimes \mathbb{1}](P_{EPR})\|_\infty, \tag{D.15}$$

where $\delta\Phi = \Phi_{\mathcal{E}} - \Phi_a^{G,\mathrm{TI}}$, and P_{EPR} is the EPR state on $\mathcal{H}^{\otimes k} \otimes \mathcal{H}^{\otimes k}$.

Proof. The proof of the lemma proceeds step-by-step in the exact same manner as the proof of Lemma 5. The additional factor of $|G|^k$ arises from the sum over symmetry operators $R_{\boldsymbol{g}}$. □

We can now analyze the symmetric translation-invariant two-layer circuit. Similar to in Appendix D.2, we have $U = U_2^{\otimes m} \cdot U_1^{\otimes m}$, where U_1, U_2 are symmetric random unitaries on 2ξ qubits each, and $U_1^{\otimes m}$ acts on even pairs of patches and $U_2^{\otimes m}$ acts on odd pairs of patches. Our main result is as follows.

Theorem 13 (Gluing symmetric translation-invariant random unitaries)**.** *For any discrete on-site symmetry group G and any $\varepsilon \leq 1$. Consider the symmetric translation-invariant two-layer brickwork circuit, U, described above. Suppose that the small random unitaries, U_1 and U_2, are each drawn from $\frac{\varepsilon}{4}$-approximate symmetric unitary mk-designs. Then U forms an ε-approximate symmetric translation-invariant unitary k-design, as long as $\xi \geq \log_2(32|G|n^6k^6/\varepsilon)$.*

By drawing each of U_1 and U_2 from a symmetric PRU ensemble on 2ξ qubits with security against sub-exponential time quantum adversaries (Appendix B), Theorem 13 immediately allows one to construct symmetric translation-invariant pseudorandom unitaries in extremely low depth.

Corollary 3 (Symmetric translation-invariant PRUs in extremely low depth)**.** *Symmetric translation-invariant PRUs exist, and can be formed in depth* $\mathrm{poly}(\log n)$ *in geometrically-local circuits, and* $\mathrm{poly}\log\log n$ *in all-to-all connected circuits.*

Proof of Theorem 13. We proceed as in the proof of Theorem 12. The twirl when both U_1 and U_2 are drawn from the symmetric Haar ensembles is approximated by,

$$\begin{aligned}&[\Phi_{\mathcal{E}} \otimes \mathbb{1}](P_{\mathrm{EPR}})\\ &= \frac{1}{D^{3k}} \sum_{\sigma,\tilde{\sigma}} \sum_{\boldsymbol{g},\tilde{\boldsymbol{g}}} \mathrm{tr}_e((\mathcal{T} R_{\tilde{\boldsymbol{g}}} \tilde{\sigma} \mathcal{T}^{-1})^{-1} R_{\boldsymbol{g}} \sigma) \cdot \mathrm{tr}_o(\tilde{\sigma}^{-1} R_{\tilde{\boldsymbol{g}}}^{-1} R_{\boldsymbol{g}} \sigma)\\ &\qquad \cdot [(\mathcal{T} R_{\tilde{\boldsymbol{g}}} \tilde{\sigma} \mathcal{T}^{-1})_e \otimes (R_{\tilde{\boldsymbol{g}}} \tilde{\sigma})_o] \otimes [(R_{\boldsymbol{g}} \sigma)_e \otimes (R_{\boldsymbol{g}} \sigma)_o],\end{aligned} \tag{D.16}$$

up to relative error $(1 + |G|K^2/D^2)^2$. The pair of traces are equal to

$$\mathrm{tr}_o(\tilde{\sigma}^{-1} R_{\tilde{\boldsymbol{g}}}^{-1} R_{\boldsymbol{g}} \sigma) = D^{-|\tilde{\sigma}^{-1}\sigma|} \cdot \prod_{\ell \in \sigma\tilde{\sigma}^{-1}} \delta\Big(\prod_{i\in\ell} \tilde{g}_i g_i^{-1} = e\Big) \tag{D.17}$$

$$\mathrm{tr}_e((\mathcal{T} R_{\tilde{\boldsymbol{g}}} \tilde{\sigma} \mathcal{T}^{-1})^{-1} R_{\boldsymbol{g}} \sigma) = D^{-|\tilde{\sigma}^{-1}\mathcal{T}^{-1}\sigma\mathcal{T}|} \cdot \prod_{\ell \in \tilde{\sigma}^{-1}\mathcal{T}^{-1}\sigma\mathcal{T}} \delta\Big(\prod_{i\in\ell} \tilde{g}_i g_i^{-1} = e\Big), \tag{D.18}$$

which is identical to the translation-invariant circuit without symmetries, except for the additional condition that the symmetry operators within each cycle ℓ product to the identity. This uniquely fixes $\tilde{\boldsymbol{g}}$ within each unit cycle. The analogous expression for the approximate symmetric translation-invariant Haar twirl is

$$[\Phi_a^{G,\mathrm{TI}} \otimes \mathbb{1}](P_{\mathrm{EPR}}) = \frac{1}{D^{2k}} \sum_{\sigma\,:\,[\sigma,\mathcal{T}^{\otimes k}]=0} \sum_{\boldsymbol{g}} [(R_{\boldsymbol{g}}\sigma)_e \otimes (R_{\boldsymbol{g}}\sigma)_o] \otimes [(R_{\boldsymbol{g}}\sigma)_e \otimes (R_{\boldsymbol{g}}\sigma)_o]. \tag{D.19}$$

Taking the difference of the two and applying the triangle inequality gives,

$$\|[\delta\Phi \otimes \mathbb{1}](P_{\mathrm{EPR}})\|_\infty \tag{D.20}$$

$$\leq \frac{1}{D^{2k}} \left(\sum_{\sigma,\tilde{\sigma}} \sum_{\boldsymbol{g},\tilde{\boldsymbol{g}}} D^{-|\tilde{\sigma}^{-1}\mathcal{T}^{-1}\sigma\mathcal{T}|} D^{-|\tilde{\sigma}^{-1}\sigma|} \prod_{\ell \in \tilde{\sigma}^{-1}\mathcal{T}^{-1}\sigma\mathcal{T}} \delta\Big(\prod_{i\in\ell} \tilde{g}_i g_i^{-1}\Big) \prod_{\ell \in \sigma\tilde{\sigma}^{-1}} \delta\Big(\prod_{i\in\ell} \tilde{g}_i g_i^{-1}\Big) - |G|^k n^k k! \right), \tag{D.21}$$

As in the proof of the symmetric gluing lemma, Lemma 4, we can upper bound the sum over $\tilde{\boldsymbol{g}}$ by noting that each cycle (of either permutation) uniquely fixes one value of $\tilde{g}_i$. There are therefore, at most, $|G|^{\min(|\tilde{\sigma}^{-1}\sigma|,|\tilde{\sigma}^{-1}\mathcal{T}^{-1}\sigma\mathcal{T}|)}$ values that $\tilde{\boldsymbol{g}}$ can take. Using $\min(x,y) \leq x + y$, we can upper bound

our previous expression by,

$$\begin{aligned}\|[\delta\Phi\otimes\mathbb{1}](P_{\mathrm{EPR}})\|_\infty &\leq \frac{1}{D^{2k}}\left(\sum_{\sigma,\tilde{\sigma}}\sum_{\boldsymbol{g}}(D/|G|)^{-|\tilde{\sigma}^{-1}\mathcal{T}^{-1}\sigma\mathcal{T}|}(D/|G|)^{-|\tilde{\sigma}^{-1}\sigma|}-|G|^k n^k k!\right)\\ &=\frac{|G|^k}{D^{2k}}\left(\sum_{\sigma,\tilde{\sigma}}(D/|G|)^{-|\tilde{\sigma}^{-1}\mathcal{T}^{-1}\sigma\mathcal{T}|}(D/|G|)^{-|\tilde{\sigma}^{-1}\sigma|}-n^k k!\right),\end{aligned} \tag{D.22}$$

where in the second line we perform the sum over $\boldsymbol{g}$, which yields a factor of $|G|^k$. The sum is identical to the sum that we bounded in the proof of Theorem 12, replacing $D \to D/|G|$. Making this replacement in the bounds in Theorem 12 completes our proof. □

E Quantum phases of matter

In this section, we provide full details on the implications of our results to the hardness of recognizing phases of matter in quantum systems. We begin by precisely stating the definitions of phases of matter that our results apply to. We then provide detailed proofs of our Theorems 1 and 3 from the main text.

There exist many closely-related definitions of phases of matter in the literature. Our results apply to any definitions that obey the following criteria.

Definition 4 (Pure state phases of matter)**.** *Our results apply to any definition of pure state phases of matter such that:*

- *For each phase, one can associate a reference "fixed point" state $|\psi_0\rangle$.*
- *Any state $|\psi\rangle$ that can be mapped to $|\psi_0\rangle$ via a symmetric geometrically-local depth-ℓ light-cone-ξ unitary circuit, $U|\psi\rangle = |\psi_0\rangle$, is in the same phase as $|\psi_0\rangle$.*

Definition 5 (Mixed state phases of matter)**.** *Our results apply to any definition of mixed state phases of matter such that:*

- *For each phase, one can associate a reference "fixed point" state ρ_0.*
- *Any state ρ that can be mapped to and from ρ_0 via symmetry-covariant geometrically-local depth-ℓ light-cone-ξ quantum channels, $\mathcal{C}_1(\rho) = \rho_0$ and $\mathcal{C}_2(\rho_0) = \rho$, is in the same phase as ρ_0.*

To be precise, in the second definition, the only channels that we require are combinations of symmetric local unitary circuits and adding and removing ancilla qubits in the maximally mixed state.

E.1 Proof of Theorem 1: Hardness of recognizing pure state phases of matter

Let us first address the setting when $\xi = \omega(\log n)$. Here, Theorem 1 follows immediately from our construction of symmetric PRUs on n qubits in Theorem 3. Consider two quantum states $|\psi_{\mathrm{P1}}\rangle = U|\psi_{0,\mathrm{P1}}\rangle$ and $|\psi_{\mathrm{P2}}\rangle = U|\psi_{0,\mathrm{P2}}\rangle$. Here, P1 and P2 denote two different phases of matter (for example, a trivial and non-trivial phase). The states $|\psi_{0,\mathrm{P1}}\rangle$ and $|\psi_{0,\mathrm{P2}}\rangle$ denote their fixed point states. The states $|\psi_{\mathrm{P1}}\rangle$ and $|\psi_{\mathrm{P2}}\rangle$ are random and are obtained by drawing U from the two-layer circuit with light-cone ξ, circuit depth $\mathrm{poly}(\xi)$, and security against any $\exp(o(\xi))$-time adversary (Theorem 3). Without loss of generality, we assume that $|\psi_{0,\mathrm{P1}}\rangle$ and $|\psi_{0,\mathrm{P2}}\rangle$ are symmetric, since any phase of matter has a symmetric representative. By the definition of a symmetric PRU, U cannot be distinguished from a symmetric Haar-random unitary by any $\exp(o(\xi))$-time algorithm. Hence, neither $|\psi_{\mathrm{P1}}\rangle$ nor

$|\psi_{\mathrm{P2}}\rangle$ can be distinguished from a Haar-random symmetric state. Hence, the two random states also cannot be distinguished from each other. If the phase of matter could be recognized, then one could use this to distinguish the states. Hence, the phase of matter cannot be recognized.

We now address smaller values of ξ, i.e. $\xi = \mathcal{O}(\log n)$ or smaller. In this regime, the required circuit size will always be polynomial or smaller in n, but may still grow exponentially in the value of ξ. As before, we draw the random unitary U from the two-layer circuit ensemble with light-cone ξ, however ξ is now too small for the two-layer circuit to form a PRU on n qubits. Let C denote the circuit size of the quantum or classical algorithm. By assumption, C is smaller than any exponential in ξ, hence C is smaller than any constant fractional power of n. Now, any circuit of size C can act on at most C qubits of the unknown quantum state. Hence, we can trace out all but C of the small random unitaries in the two-layer circuit. The output of the experiment is unaffected by the details of the two-layer circuit in the traced out regions. Thus, the output would be the same if the $\mathcal{O}(n)$ traced out small random unitaries were replaced with $\mathcal{O}(C)$ small random unitaries that glue together the remaining un-traced out regions that the algorithm acts on. Hence, the output is identical to the output of a fictitious two-layer circuit that acts on only $\mathcal{O}(C)$ qubits instead of n qubits. The fictitious two-layer circuit forms a symmetric PRU with security against any poly(C)-time adversary whenever $\xi = \omega(\log C)$. Hence, recognizing the phase of matter requires $\xi = \mathcal{O}(\log C)$, which implies that the circuit size must grow exponentially in ξ. □

E.2 Proof of Theorem 5: Hardness of recognizing mixed state phases of matter

The proof follows by an identical argument to Theorem 1 with only small modifications. Let ρ_0 denote the fixed point of any mixed-state phase of matter. Our first step is to take the tensor product of ρ_0 with n ancilla qudits in the maximally mixed state, yielding $\tilde{\rho}_0 \equiv \rho_0 \otimes (\mathbb{1}/d)^{\otimes n}$ where d is the local qudit dimension. Here, the system symmetry G does not act on the ancilla qubits. We assume that each ancilla qudit is placed geometrically next to a corresponding qudit in the original system. Our second step is to apply a two-layer circuit of small symmetric PRUs on 2ξ pairs of qudits each. This yields a random state $\rho = U\tilde{\rho}_0 U^\dagger$ in the same phase as $\tilde{\rho}_0$ and hence ρ_0.

We will now show that the random state ρ cannot be distinguished a mixture of maximally mixed states in each symmetry sector. Let $\rho_{\mathrm{mixed}} = \bigotimes_\lambda p_\lambda(\mathbb{1}_{\mathcal{M}_\lambda} \otimes \rho_{\mathcal{F}_\lambda})$ denote the mixed state obtained by completely depolarizing ρ (or equivalently, $\tilde{\rho}_0$) in each symmetry sector. Here, the factor of p_λ is chosen such that $\mathrm{tr}(\rho_{\mathcal{F}_\lambda}) = 1$. We assume $\xi = \omega(\log n)$; the proof for smaller ξ follows by an identical argument as in the proof of Theorem 1. Since U is a symmetric PRU, we know that ρ cannot be distinguished from the $V\tilde{\rho}_0 V^\dagger$ where V is a symmetric Haar-random unitary. It remains only to show that $V\tilde{\rho}_0 V^\dagger$ is indistinguishable from ρ_{mixed}.

To show this, we borrow a technique from recent studies of pseudorandom states and unitaries known as the distinct subspace. We refer to Refs. [32, 44] for a thorough introduction. Let k denote the number of queries that an algorithm makes to the random state of interest. Let $x^{(j)}$ label a computational basis of d^n bitstrings for the n ancilla qudits in the j-th copy of state, where $j = 1, \ldots, k$. Then, let Π^{dist} denote the projector onto the subspace of the k copies of n ancilla qudits where all k values of $x^{(j)}$ are distinct. The projector is upper bounded by the operator inequality, $\mathbb{1} - \Pi^{\mathrm{dist}} \preceq \sum_{i<j} \Pi^{\mathrm{eq},ij}$, where $\Pi^{\mathrm{eq},ij}$ is the projector onto states where $x^{(i)} = x^{(j)}$. The first step of our proof is then to project the k copies of the state $\tilde{\rho}_0$ onto the distinct subspace on the ancilla qudits. This incurs a negligibly small error,

$$\left\| \tilde{\rho}_0^{\otimes k} - \Pi^{\mathrm{dist}} \tilde{\rho}_0^{\otimes k} \Pi^{\mathrm{dist}} \right\|_1 = \mathrm{tr}\Big((\mathbb{1} - \Pi^{\mathrm{dist}}) \tilde{\rho}_0^{\otimes k} \Big) \leq \sum_{i<j} \mathrm{tr}\Big(\Pi^{\mathrm{eq},ij} \tilde{\rho}_0^{\otimes k} \Big) \leq k^2/2d^n. \tag{E.1}$$

The term is upper bounded by $k^2/2d^n$, since the probability for two copies of the n ancilla qudits to be in the same state is $1/d^n$. The factor of $k(k-1)/2 \leq k^2/2$ counts the number of terms in the sum.

The second step of our proof is to compute the twirl of the projected state over the symmetric Haar-random unitary V. From Lemma 2, the twirl is equal to the following up to negligibly small relative error,

$$\frac{1}{D^k}\sum_{\pi}\sum_{\boldsymbol{g}}\mathrm{tr}\Big(\Pi^{\mathrm{dist}}\tilde{\rho}_0^{\otimes k}\Pi^{\mathrm{dist}}\pi^{-1}R_{\boldsymbol{g}}^{-1}\Big)R_{\boldsymbol{g}}\pi=\frac{1}{D^k}\sum_{\boldsymbol{g}}\mathrm{tr}\Big(\Pi^{\mathrm{dist}}\tilde{\rho}_0^{\otimes k}\Pi^{\mathrm{dist}}R_{\boldsymbol{g}}^{-1}\Big)R_{\boldsymbol{g}}, \tag{E.2}$$

where we use that $\mathrm{tr}_{\mathrm{ancillas}}(\Pi^{\mathrm{dist}}\pi^{-1})=\delta_{\mathbb{1},\pi}$. From Eq. (E.1), this state is equal to

$$\frac{1}{D^k}\sum_{\boldsymbol{g}}\mathrm{tr}\Big(\tilde{\rho}_0^{\otimes k}R_{\boldsymbol{g}}^{-1}\Big)R_{\boldsymbol{g}}=\rho_{\mathrm{mixed}}^{\otimes k} \tag{E.3}$$

up to an additional 1-norm error $k^2/2d$. This completes the proof. □

F Classical phases of matter

In this section, we provide the full details of our results on the hardness of recognizing phases of matter in classical systems. We begin by precisely stating our definition of classical phases of matter, which is closely modeled off of standard definitions for quantum phases of matter (restricted to classical states and channels). We then provide the detailed proof of our Theorems 6 from the main text.

To our knowledge, a precise definition of classical phases of matter, in a similar format to the now standard definitions of quantum phases of matter, has not been given in the literature. We propose the following definition, which is identical to the definition of mixed state quantum phases of matter but with all operations restricted to be entirely classical.

Definition 6 (Classical phases of matter)**.** *Our results apply to any definition of classical phases of matter such that:*

- *For each phase, one can associate a reference "fixed point" probability distribution $p_0(x)$.*
- *Any distribution $p(x)$ that can be mapped to and from $p_0(x)$ via symmetry-covariant geometrically-local depth-ℓ stochastic processes, $\mathcal{M}_1(p)=p_0$ and $\mathcal{M}_2(p_0)=p$, is in the same phase as $p_0(x)$.*

As in the mixed state quantum setting, the only stochastic processes that we require are combinations of symmetric local reversible circuits and adding and removing ancilla bits in the maximally mixed state.

F.1 Proof of Theorem 6: Hardness of recognizing classical phases of matter

Let ρ_0 denote the diagonal density matrix with entries $\langle x|\,\rho_0\,|x\rangle = p_0(x)$. Let $\mathcal{M}_P(\rho)=P[\rho\otimes(\mathbb{1}/2)^{\otimes n}]P^\dagger$ denote the classical stochastic process of interest [Fig. 4(a)], where $P=\bigotimes_{\alpha=1}^m P_\alpha$ and each P_α is a pseudorandom permutation acting on $2\xi=2n/m$ bits. By assumption, each pseudorandom permutation cannot be distinguished from a truly random permutation by any poly n-time classical or quantum experiment. Hence, from hereon, we can assume each P_α is truly random.

Consider a classical or quantum experiment that queries $\mathcal{M}(p_0(x))$ at most $k=\mathrm{poly}\,n$ times and outputs 0 if it decides that the state is in the trivial phase and 1 if it decides the state is in the symmetry-breaking phase. The expected output of the experiment can be written as,

$$\mathop{\mathbb{E}}_{P}\Big[\mathrm{tr}\Big(M\cdot\mathcal{M}_P(\rho_0)^{\otimes k}\Big)\Big], \tag{F.1}$$

for some Hermitian observable M acting on all k copies, with $0\preceq M\preceq 1$.

To proceed, let us define the *symmetric distinct subspace* as the set of all sets of k bitstrings such that no two strings are equal up to the symmetry operation. That is, $x_i \neq x_j$ for all $1 \leq i < j \leq k$ and $x_i \neq \bar{x}_j$ as well, where $\bar{x}_j$ is equal to x_j with all bits flipped (i.e. acted on by the symmetry operation). We let $\Pi_\alpha^{\text{sym-dist}}$ denote the projector onto the symmetric distinct subspace on patch α; that is, onto all sets of k 2ξ-bit strings on patch i such that no two strings are equal to one another up to the symmetry. We also let $\Pi_*^{\text{sym-dist}} = \bigotimes_{\alpha=1}^m \Pi_\alpha^{\text{dist}}$ denote the projector onto the distinct subspace on every patch.

The first step of our proof is to replace $\rho_0 \otimes (\mathbb{1}/2)^{\otimes n}$ with its projection onto the symmetric distinct subspace $\Pi_*^{\text{sym-dist}}$. This incurs an error

$$E \equiv \left\| \rho_0 \otimes (\mathbb{1}/2)^{\otimes n} - \Pi_*^{\text{sym-dist}} \left[\rho_0 \otimes (\mathbb{1}/2)^{\otimes n} \right] \Pi_*^{\text{sym-dist}} \right\|_1 = \text{tr}\left(\left(\mathbb{1} - \Pi_*^{\text{sym-dist}} \right) \left[\rho_0 \otimes (\mathbb{1}/2)^{\otimes n} \right] \right), \tag{F.2}$$

where we use that $\Pi_*^{\text{sym-dist}}$ and the state commute since the state is classical. The trace measures the probability that any non-distinct set of k 2ξ-bit strings is observed when measuring $\rho_0 \otimes (\mathbb{1}/2)^{\otimes n}$. This can be upper bounded using the inequality $\Pi_*^{\text{sym-dist}} \preceq \sum_\alpha \sum_{i<j} \Pi_\alpha^{\text{eq},ij}$, where $\Pi_\alpha^{\text{eq},ij}$ is the projector onto all states with the same bitstring (up to the symmetry operation) on copies i and j. Here, the first sum is over all patches and the second sum is over all pairs of copies, $1 \leq i < j \leq k$. The probability to observe the same (up to the symmetry) 2ξ-bit string on copies i and j is upper bounded by the probability to observe the same (up to the symmetry) ξ-bit string when restricting attention to the ξ ancilla qubits. Since the ancilla qubits are maximally mixed, this probability is equal to $2/2^\xi$, where the factor of 2 in the numerator accounts for equality up to the symmetry operation. From this, the 1-norm in Eq. (F.2) is upper bounded by $mk(k-1)/2^\xi$.

To complete the proof, we note that the action of each symmetric random permutation on a symmetry-distinct set of bitstrings is especially simple,

$$\mathop{\mathbb{E}}_{P_\alpha} \left[\, |x_\alpha\rangle\langle x_\alpha| \, \right] = \Pi_\alpha^{\text{sym-dist}} / \mathcal{D}_\alpha^{\text{sym-dist}}, \quad \text{if} \quad \Pi_\alpha^{\text{sym-dist}} |x_\alpha\rangle = |x_\alpha\rangle, \tag{F.3}$$

where $\mathcal{D}_\alpha^{\text{sym-dist}} = 2^{2\xi}(2^{2\xi}-2)(2^{2\xi}-4)\cdots(2^{2\xi}-2k+2) \geq 2^{2\xi}(1-k^2/2^{2\xi})$ is the dimension of the symmetric distinct subspace on patch α. The formula follows because a symmetric random permutation sends each input bitstring to a random output bitstring, up to only the condition that two inputs related by the symmetry operation must go to outputs that are also related by the symmetry. Since x_α is symmetry-distinct, no two inputs are related by the symmetry. Hence, all outputs are distinct, not related by the symmetry, and otherwise independently random. When taken in expectation, this produces the maximally mixed state on the symmetric distinct subspace.

Applying Eq. (F.3) to each patch α, we find

$$\mathop{\mathbb{E}}_{P} \left[\left(P \left[\Pi_*^{\text{sym-dist}} \left[\rho_0 \otimes (\mathbb{1}/2)^{\otimes n} \right] \Pi_*^{\text{sym-dist}} \right] P^\dagger \right)^{\otimes k} \right] = (1-E) \cdot \bigotimes_{\alpha=1}^m \left(\frac{\Pi_\alpha^{\text{sym-dist}}}{\mathcal{D}_\alpha^{\text{sym-dist}}} \right), \tag{F.4}$$

where $1-E$ is the normalization of the state, with E defined in Eq. (F.2). The state on the right hand side is close to the fixed state,

$$\bigotimes_{\alpha=1}^m \left(\frac{\Pi_\alpha^{\text{sym-dist}}}{\mathcal{D}_\alpha^{\text{sym-dist}}} \right), \tag{F.5}$$

up to 1-norm error $E \leq mk(k-1)/2^\xi$. However, this state is independent of the input probability distribution $p_0(x)$. Hence, the expected output states of any two different input distributions, $p_0(x)$ and $p_1(x)$, will be equal up to 1-norm error $4mk(k-1)/2^\xi$. This immediately implies that their expected output decisions are equal up to $4mk(k-1)/2^\xi = 1/\omega(\text{poly}\, n)$ given $\xi = \omega(\log n)$. Hence, the experiment cannot distinguish the output distributions arising from $p_0(x)$ and $p_1(x)$ in polynomial time. Setting $p_1(x)$ to be the maximally mixed distribution yields Theorem 6. □

References

[1] Lev Davidovich Landau and Evgenii Mikhailovich Lifshitz. *Statistical Physics: Volume 5*, volume 5. Elsevier, 2013.

[2] Xiao-Gang Wen. *Quantum field theory of many-body systems: From the origin of sound to an origin of light and electrons*. Oxford university press, 2004.

[3] Xie Chen, Zheng-Cheng Gu, Zheng-Xin Liu, and Xiao-Gang Wen. Symmetry protected topological orders and the group cohomology of their symmetry group. *Physical Review B—Condensed Matter and Materials Physics*, 87(15):155114, 2013.

[4] Bei Zeng, Xie Chen, Duan-Lu Zhou, Xiao-Gang Wen, et al. *Quantum information meets quantum matter*. Springer, 2019.

[5] Ehud Altman, Kenneth R Brown, Giuseppe Carleo, Lincoln D Carr, Eugene Demler, Cheng Chin, Brian DeMarco, Sophia E Economou, Mark A Eriksson, Kai-Mei C Fu, et al. Quantum simulators: Architectures and opportunities. *PRX quantum*, 2(1):017003, 2021.

[6] Bela Bauer, Sergey Bravyi, Mario Motta, and Garnet Kin-Lic Chan. Quantum algorithms for quantum chemistry and quantum materials science. *Chemical reviews*, 120(22):12685–12717, 2020.

[7] Giulia Semeghini, Harry Levine, Alexander Keesling, Sepehr Ebadi, Tout T Wang, Dolev Bluvstein, Ruben Verresen, Hannes Pichler, Marcin Kalinowski, Rhine Samajdar, et al. Probing topological spin liquids on a programmable quantum simulator. *Science*, 374(6572):1242–1247, 2021.

[8] Hsin-Yuan Huang, Richard Kueng, Giacomo Torlai, Victor V Albert, and John Preskill. Provably efficient machine learning for quantum many-body problems. *Science*, 377(6613):eabk3333, 2022.

[9] Logan W Clark, Nathan Schine, Claire Baum, Ningyuan Jia, and Jonathan Simon. Observation of laughlin states made of light. *Nature*, 582(7810):41–45, 2020.

[10] KJ Satzinger, Y-J Liu, A Smith, C Knapp, M Newman, C Jones, Z Chen, C Quintana, X Mi, A Dunsworth, et al. Realizing topologically ordered states on a quantum processor. *Science*, 374(6572):1237–1241, 2021.

[11] Xu Zhang, Wenjie Jiang, Jinfeng Deng, Ke Wang, Jiachen Chen, Pengfei Zhang, Wenhui Ren, Hang Dong, Shibo Xu, Yu Gao, et al. Digital quantum simulation of floquet symmetry-protected topological phases. *Nature*, 607(7919):468–473, 2022.

[12] Maxime Dupont and Joel E Moore. Quantum criticality using a superconducting quantum processor. *Physical Review B*, 106(4):L041109, 2022.

[13] Mohsin Iqbal, Nathanan Tantivasadakarn, Ruben Verresen, Sara L Campbell, Joan M Dreiling, Caroline Figgatt, John P Gaebler, Jacob Johansen, Michael Mills, Steven A Moses, et al. Non-abelian topological order and anyons on a trapped-ion processor. *Nature*, 626(7999):505–511, 2024.

[14] Julian Léonard, Sooshin Kim, Joyce Kwan, Perrin Segura, Fabian Grusdt, Cécile Repellin, Nathan Goldman, and Markus Greiner. Realization of a fractional quantum hall state with ultracold atoms. *Nature*, 619(7970):495–499, 2023.

[15] Reza Haghshenas, Eli Chertkov, Michael Mills, Wilhelm Kadow, Sheng-Hsuan Lin, Yi-Hsiang Chen, Chris Cade, Ido Niesen, Tomislav Begušić, Manuel S Rudolph, et al. Digital quantum magnetism at the frontier of classical simulations. *arXiv preprint arXiv:2503.20870*, 2025.

[16] Melissa Will, TA Cochran, Eliott Rosenberg, Bernhard Jobst, Norhan Mahmoud Eassa, P Roushan, M Knap, A Gammon-Smith, and F Pollmann. Probing non-equilibrium topological order on a quantum processor. *arXiv preprint arXiv:2501.18461*, 2025.

[17] Hong-Chen Jiang, Zhenghan Wang, and Leon Balents. Identifying topological order by entanglement entropy. *Nature Physics*, 8(12):902–905, 2012.

[18] Joaquin F Rodriguez-Nieva and Mathias S Scheurer. Identifying topological order through unsupervised machine learning. *Nature Physics*, 15(8):790–795, 2019.

[19] Ze-Pei Cian, Hossein Dehghani, Andreas Elben, Benoît Vermersch, Guanyu Zhu, Maissam Barkeshli, Peter Zoller, and Mohammad Hafezi. Many-body chern number from statistical correlations of randomized measurements. *Physical Review Letters*, 126(5):050501, 2021.

[20] Johannes Herrmann, Sergi Masot Llima, Ants Remm, Petr Zapletal, Nathan A McMahon, Colin Scarato, François Swiadek, Christian Kraglund Andersen, Christoph Hellings, Sebastian Krinner, et al. Realizing quantum convolutional neural networks on a superconducting quantum processor to recognize quantum phases. *Nature communications*, 13(1):4144, 2022.

[21] Hsin-Yuan Huang, Richard Kueng, Giacomo Torlai, Victor V Albert, and John Preskill. Provably efficient machine learning for quantum many-body problems. *Science*, 377(6613):eabk3333, 2022.

[22] Ze-Pei Cian, Mohammad Hafezi, and Maissam Barkeshli. Extracting wilson loop operators and fractional statistics from a single bulk ground state. *arXiv preprint arXiv:2209.14302*, 2022.

[23] Iris Cong, Soonwon Choi, and Mikhail D Lukin. Quantum convolutional neural networks. *Nature Physics*, 15(12):1273–1278, 2019.

[24] Ethan Lake, Shankar Balasubramanian, and Soonwon Choi. Exact quantum algorithms for quantum phase recognition: Renormalization group and error correction. *arXiv preprint arXiv:2211.09803*, 2022.

[25] Iris Cong, Nishad Maskara, Minh C Tran, Hannes Pichler, Giulia Semeghini, Susanne F Yelin, Soonwon Choi, and Mikhail D Lukin. Enhancing detection of topological order by local error correction. *Nature Communications*, 15(1):1527, 2024.

[26] Adam Bouland, Bill Fefferman, Soumik Ghosh, Tony Metger, Umesh Vazirani, Chenyi Zhang, and Zixin Zhou. Public-key pseudoentanglement and the hardness of learning ground state entanglement structure. *arXiv preprint arXiv:2311.12017*, 2023.

[27] Hsin-Yuan Huang, Yunchao Liu, Michael Broughton, Isaac Kim, Anurag Anshu, Zeph Landau, and Jarrod R McClean. Learning shallow quantum circuits. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, pages 1343–1351, 2024.

[28] Hyun-Soo Kim, Isaac H Kim, and Daniel Ranard. Learning state preparation circuits for quantum phases of matter. *arXiv preprint arXiv:2410.23544*, 2024.

[29] Zeph Landau and Yunchao Liu. Learning quantum states prepared by shallow circuits in polynomial time. In *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*, pages 1828–1838, 2025.

[30] Thomas Schuster, Jonas Haferkamp, and Hsin-Yuan Huang. Random unitaries in extremely low depth. *arXiv preprint arXiv:2407.07754*, 2024.

[31] Zhengfeng Ji, Yi-Kai Liu, and Fang Song. Pseudorandom quantum states. In *Advances in Cryptology–CRYPTO 2018: 38th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 19–23, 2018, Proceedings, Part III 38*, pages 126–152. Springer, 2018.

[32] Tony Metger, Alexander Poremba, Makrand Sinha, and Henry Yuen. Simple constructions of linear-depth t-designs and pseudorandom unitaries. *arXiv preprint arXiv:2404.12647*, 2024.

[33] Chi-Fang Chen, Adam Bouland, Fernando GSL Brandão, Jordan Docter, Patrick Hayden, and Michelle Xu. Efficient unitary designs and pseudorandom unitaries from permutations. *arXiv preprint arXiv:2404.16751*, 2024.

[34] Fermi Ma and Hsin-Yuan Huang. How to construct random unitaries. *arXiv preprint arXiv:2410.10116*, 2024.

[35] Thomas Schuster, Fermi Ma, Alex Lombardi, Fernando Brandao, and Hsin-Yuan Huang. Strong random unitaries and fast scrambling. *arXiv preprint arXiv:2509.26310*, 2025.

[36] Joseph Emerson, Yaakov S Weinstein, Marcos Saraceno, Seth Lloyd, and David G Cory. Pseudo-random unitary operators for quantum information processing. *science*, 302(5653):2098–2100, 2003.

[37] David Gross, Koenraad Audenaert, and Jens Eisert. Evenly distributed unitaries: On the structure of unitary designs. *Journal of mathematical physics*, 48(5), 2007.

[38] Christoph Dankert. Efficient simulation of random quantum states and operators. *arXiv preprint quant-ph/0512217*, 2005.

[39] Christoph Dankert, Richard Cleve, Joseph Emerson, and Etera Livine. Exact and approximate unitary 2-designs and their application to fidelity estimation. *Physical Review A*, 80(1):012304, 2009.

[40] Fernando GSL Brandao, Aram W Harrow, and Michał Horodecki. Local random quantum circuits are approximate polynomial-designs. *Communications in Mathematical Physics*, 346:397–434, 2016.

[41] Jeongwan Haah, Yunchao Liu, and Xinyu Tan. Efficient approximate unitary designs from random pauli rotations. *arXiv preprint arXiv:2402.05239*, 2024.

[42] Chi-Fang Chen, Jeongwan Haah, Jonas Haferkamp, Yunchao Liu, Tony Metger, and Xinyu Tan. Incompressibility and spectral gaps of random circuits. *arXiv preprint arXiv:2406.07478*, 2024.

[43] Nicholas LaRacuente and Felix Leditzky. Approximate unitary k-designs from shallow, low-communication circuits. *arXiv preprint arXiv:2407.07876*, 2024.

[44] Laura Cui, Thomas Schuster, Fernando Brandao, and Hsin-Yuan Huang. Unitary designs in nearly optimal depth. *arXiv preprint arXiv:2507.06216*, 2025.

[45] Yuzhen Zhang, Sagar Vijay, Yingfei Gu, and Yimu Bao. Designs from magic-augmented clifford circuits. *arXiv preprint arXiv:2507.02828*, 2025.

[46] Prabhanjan Ananth, Luowen Qian, and Henry Yuen. Cryptography from pseudorandom quantum states. In *Annual International Cryptology Conference*, pages 208–236. Springer, 2022.

[47] William Kretschmer, Luowen Qian, Makrand Sinha, and Avishay Tal. Quantum cryptography in algorithmica. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, pages 1589–1602, 2023.

[48] Joseph Emerson, Robert Alicki, and Karol Życzkowski. Scalable noise estimation with random unitary operators. *Journal of Optics B: Quantum and Semiclassical Optics*, 7(10):S347, 2005.

[49] Emanuel Knill, Dietrich Leibfried, Rolf Reichle, Joe Britton, R Brad Blakestad, John D Jost, Chris Langer, Roee Ozeri, Signe Seidelin, and David J Wineland. Randomized benchmarking of quantum gates. *Physical Review A*, 77(1):012307, 2008.

[50] Andreas Elben, Steven T Flammia, Hsin-Yuan Huang, Richard Kueng, John Preskill, Benoît Vermersch, and Peter Zoller. The randomized measurement toolbox. *Nature Reviews Physics*, 5(1):9–24, 2023.

[51] Frank Arute, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph C Bardin, Rami Barends, Rupak Biswas, Sergio Boixo, Fernando GSL Brandao, David A Buell, et al. Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779):505–510, 2019.

[52] Alexis Morvan, B Villalonga, X Mi, S Mandra, A Bengtsson, PV Klimov, Z Chen, S Hong, C Erickson, IK Drozdov, et al. Phase transition in random circuit sampling. *arXiv preprint arXiv:2304.11119*, 2023.

[53] Dmitry A Abanin, Rajeev Acharya, Laleh Aghababaie-Beni, Georg Aigeldinger, Ashok Ajoy, Ross Alcaraz, Igor Aleiner, Trond I Andersen, Markus Ansmann, Frank Arute, et al. Constructive interference at the edge of quantum ergodic dynamics. *arXiv preprint arXiv:2506.10191*, 2025.

[54] Hsin-Yuan Huang, Michael Broughton, Jordan Cotler, Sitan Chen, Jerry Li, Masoud Mohseni, Hartmut Neven, Ryan Babbush, Richard Kueng, John Preskill, et al. Quantum advantage in learning from experiments. *Science*, 376(6598):1182–1186, 2022.

[55] Sitan Chen, Jordan Cotler, Hsin-Yuan Huang, and Jerry Li. Exponential separations between learning with and without quantum memory. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 574–585. IEEE, 2022.

[56] Sitan Chen, Jordan Cotler, Hsin-Yuan Huang, and Jerry Li. A hierarchy for replica quantum advantage. *arXiv:2111.05874*, 2021.

[57] Dorit Aharonov, Jordan Cotler, and Xiao-Liang Qi. Quantum algorithmic measurement. *Nature communications*, 13(1):1–9, 2022.

[58] Jordan Cotler, Thomas Schuster, and Masoud Mohseni. Information-theoretic hardness of out-of-time-order correlators. *Physical Review A*, 108(6):062608, 2023.

[59] Yasuhiro Sekino and Leonard Susskind. Fast scramblers. *Journal of High Energy Physics*, 2008(10):065, 2008.

[60] Patrick Hayden and John Preskill. Black holes as mirrors: quantum information in random subsystems. *JHEP*, 2007(09):120, 2007.

[61] Beni Yoshida and Alexei Kitaev. Efficient decoding for the hayden-preskill protocol. *arXiv preprint arXiv:1710.03363*, 2017.

[62] Adam R Brown, Hrant Gharibyan, Stefan Leichenauer, Henry W Lin, Sepehr Nezami, Grant Salton, Leonard Susskind, Brian Swingle, and Michael Walter. Quantum gravity in the lab. i. teleportation by size and traversable wormholes. *PRX quantum*, 4(1):010320, 2023.

[63] Thomas Schuster, Bryce Kobrin, Ping Gao, Iris Cong, Emil T Khabiboulline, Norbert M Linke, Mikhail D Lukin, Christopher Monroe, Beni Yoshida, and Norman Y Yao. Many-body quantum teleportation via operator spreading in the traversable wormhole protocol. *Physical Review X*, 12(3):031013, 2022.

[64] J. M. Deutsch. Quantum statistical mechanics in a closed system. *Phys. Rev. A*, 43:2046, 1991.

[65] Mark Srednicki. Chaos and quantum thermalization. *Phys. Rev. E*, 50:888, 1994.

[66] Marcos Rigol, Vanja Dunjko, and Maxim Olshanii. Thermalization and its mechanism for generic isolated quantum systems. *Nature*, 452(7189):854–858, 2008.

[67] Matthew PA Fisher, Vedika Khemani, Adam Nahum, and Sagar Vijay. Random quantum circuits. *Annual Review of Condensed Matter Physics*, 14:335–379, 2023.

[68] Adam Nahum, Jonathan Ruhman, Sagar Vijay, and Jeongwan Haah. Quantum Entanglement Growth Under Random Unitary Dynamics. *Phys. Rev. X*, 7:031016, 2017.

[69] Jordan Cotler, Nicholas Hunter-Jones, and Daniel Ranard. Fluctuations of subsystem entropies at late times. *Physical Review A*, 105(2):022416, 2022.

[70] Linghang Kong and Zi-Wen Liu. Charge-conserving unitaries typically generate optimal covariant quantum error-correcting codes. *arXiv preprint arXiv:2102.11835*, 2021.

[71] Linghang Kong and Zi-Wen Liu. Near-optimal covariant quantum error-correcting codes from random unitaries with symmetries. *PRX Quantum*, 3(2):020314, 2022.

[72] Yosuke Mitsuhashi and Nobuyuki Yoshioka. Clifford group and unitary designs under symmetry. *PRX Quantum*, 4(4):040331, 2023.

[73] Zimu Li, Han Zheng, and Zi-Wen Liu. Efficient quantum pseudorandomness under conservation laws. *arXiv preprint arXiv:2411.04893*, 2024.

[74] Zimu Li, Han Zheng, Junyu Liu, Liang Jiang, and Zi-Wen Liu. Designs from local random quantum circuits with su (d) symmetry. *PRX Quantum*, 5(4):040349, 2024.

[75] Hanqing Liu, Austin Hulse, and Iman Marvian. Unitary designs from random symmetric quantum circuits. *arXiv preprint arXiv:2408.14463*, 2024.

[76] Yosuke Mitsuhashi, Ryotaro Suzuki, Tomohiro Soejima, and Nobuyuki Yoshioka. Unitary designs of symmetric local random circuits. *Physical Review Letters*, 134(18):180404, 2025.

[77] Jeongwan Haah. Short remarks on shallow unitary circuits. *arXiv preprint arXiv:2504.14005*, 2025.

[78] Maxwell West, Diego García-Martín, NL Diaz, M Cerezo, and Martin Larocca. No-go theorems for sublinear-depth group designs. *arXiv preprint arXiv:2506.16005*, 2025.

[79] Lorenzo Grevink, Jonas Haferkamp, Markus Heinrich, Jonas Helsen, Marcel Hinsche, Thomas Schuster, and Zoltán Zimborás. Will it glue? on short-depth designs beyond the unitary group. *arXiv preprint arXiv:2506.23925*, 2025.

[80] Xiao-Gang Wen. Colloquium: Zoo of quantum-topological phases of matter. *Reviews of Modern Physics*, 89(4):041004, 2017.

[81] Matthew B Hastings and Xiao-Gang Wen. Quasiadiabatic continuation of quantum states: The stability of topological ground-state degeneracy and emergent gauge invariance. *Phys. Rev. B*, 72(4):045141, 2005.

[82] Sergey Bravyi, Matthew B Hastings, and Frank Verstraete. Lieb-robinson bounds and the generation of correlations and topological quantum order. *Physical review letters*, 97(5):050401, 2006.

[83] Xie Chen, Zheng-Cheng Gu, and Xiao-Gang Wen. Local unitary transformation, long-range quantum entanglement, wave function renormalization, and topological order. *Physical Review B—Condensed Matter and Materials Physics*, 82(15):155138, 2010.

[84] Norbert Schuch, David Pérez-García, and Ignacio Cirac. Classifying quantum phases using matrix product states and projected entangled pair states. *Physical Review B—Condensed Matter and Materials Physics*, 84(16):165139, 2011.

[85] Xie Chen, Zheng-Cheng Gu, and Xiao-Gang Wen. Complete classification of one-dimensional gapped quantum phases in interacting spin systems. *Physical Review B—Condensed Matter and Materials Physics*, 84(23):235128, 2011.

[86] Matthew B Hastings. Locality in quantum systems. *Quantum Theory from Small to Large Scales*, 95:171–212, 2010.

[87] Sergey Bravyi, Matthew B Hastings, and Spyridon Michalakis. Topological quantum order: stability under local perturbations. *Journal of mathematical physics*, 51(9), 2010.

[88] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *Journal of the ACM (JACM)*, 56(6):1–40, 2009.

[89] Ewin Tang and John Wright. Are controlled unitaries helpful? *arXiv preprint arXiv:2508.00055*, 2025.

[90] Each qudit has Hilbert space dimension $d = \mathcal{O}(1)$. The inclusion of higher-dimensional qudits $d > 2$ is natural to incorporate symmetry groups whose size is not a power of two.

[91] Alexei Kitaev and John Preskill. Topological entanglement entropy. *Phys. Rev. Lett.*, 96(11):110404, 2006.

[92] Isaac H Kim, Michael Levin, Ting-Chun Lin, Daniel Ranard, and Bowen Shi. Universal lower bound on topological entanglement entropy. *Physical Review Letters*, 131(16):166601, 2023.

[93] Dominic J Williamson, Arpit Dua, and Meng Cheng. Spurious topological entanglement entropy from subsystem symmetries. *Physical review letters*, 122(14):140506, 2019.

[94] Kohtaro Kato and Fernando GSL Brandão. Toy model of boundary states with spurious topological entanglement entropy. *Physical Review Research*, 2(3):032005, 2020.

[95] To be specific, let us consider a 1D system with symmetry group $G = \mathbb{Z}_2 \times \mathbb{Z}_2$. The group features four symmetry sectors, which are indexed by their charge under each of the two $\mathbb{Z}_2$ generators. In the symmetry-breaking phase, the ground state manifold contains one state in each symmetry sector (four states in total). Meanwhile, in the SPT phase with open boundary conditions, the ground state manifold also contains one state in each symmetry sector. Traditionally, the two phases are distinguished because the interplay of the symmetry group's action with spatial locality is very different between the two ground state manifolds. Nonetheless, our results on low-depth symmetric PRUs imply that all aspects of spatial locality are obscured to an observer without knowledge of U. Hence, no state in the symmetry-breaking manifold can be distinguished from its corresponding state in the SPT manifold.

[96] We thank Jeongwan Haah for bringing these counter-examples to our attention.

[97] Jong Yeon Lee, Chao-Ming Jian, and Cenke Xu. Quantum criticality under decoherence or weak measurement. *PRX quantum*, 4(3):030317, 2023.

[98] Yimu Bao, Ruihua Fan, Ashvin Vishwanath, and Ehud Altman. Mixed-state topological order and the errorfield double formulation of decoherence-induced transitions. *arXiv preprint arXiv:2301.05687*, 2023.

[99] Zijian Wang, Zhengzhi Wu, and Zhong Wang. Intrinsic mixed-state topological order. *PRX Quantum*, 6(1):010314, 2025.

[100] Shengqi Sang, Yijian Zou, and Timothy H Hsieh. Mixed-state quantum phases: Renormalization and quantum error correction. *Physical Review X*, 14(3):031044, 2024.

[101] Shengqi Sang and Timothy H Hsieh. Stability of mixed-state quantum phases via finite markov length. *Physical Review Letters*, 134(7):070403, 2025.

[102] Andrea Coser and David Pérez-García. Classification of phases for mixed states via fast dissipative evolution. *Quantum*, 3:174, 2019.

[103] Hanyu Xue, Jong Yeon Lee, and Yimu Bao. Tensor network formulation of symmetry protected topological phases in mixed states. *arXiv preprint arXiv:2403.17069*, 2024.

[104] Ruochen Ma and Alex Turzillo. Symmetry-protected topological phases of mixed states in the doubled space. *PRX Quantum*, 6(1):010348, 2025.

[105] Tai-Hsuan Yang, Bowen Shi, and Jong Yeon Lee. Topological mixed states: Axiomatic approaches and phases of matter. *arXiv preprint arXiv:2506.04221*, 2025.

[106] Xiaozhou Feng, Zihan Cheng, and Matteo Ippoliti. Hardness of observing strong-to-weak symmetry breaking. *arXiv preprint arXiv:2504.12233*, 2025.

[107] Mark Zhandry. A note on quantum-secure prps. *arXiv preprint arXiv:1611.05564*, 2016.

[108] Subir Sachdev. Quantum phase transitions. *Physics world*, 12(4):33, 1999.

[109] Jarrod R McClean, Sergio Boixo, Vadim N Smelyanskiy, Ryan Babbush, and Hartmut Neven. Barren plateaus in quantum neural network training landscapes. *Nat. Commun.*, 9(1):1–6, 2018.

[110] Hsin-Yuan Huang, Richard Kueng, and John Preskill. Predicting many properties of a quantum system from very few measurements. *Nature Physics*, 16(10):1050–1057, 2020.

[111] Mildred S Dresselhaus, Gene Dresselhaus, and Ado Jorio. *Group theory: application to the physics of condensed matter*. Springer Science & Business Media, 2007.

[112] Aram W Harrow and Saeed Mehraban. Approximate unitary t-designs by short random quantum circuits using nearest-neighbor and long-range gates. *Communications in Mathematical Physics*, pages 1–96, 2023.

[113] Mark Zhandry. How to construct quantum random functions. *J. ACM*, 68(5), aug 2021.