

©Copyright 2026

Jeremy Hartse

Quantum Simulation Informed by Physical Structure

Jeremy Hartse

A dissertation
submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

University of Washington

2026

Reading Committee:

Martin Savage, Chair

Lukasz Fidkowski

Silas Beane

Program Authorized to Offer Degree:
Physics

University of Washington

Abstract

Quantum Simulation Informed by Physical Structure

Jeremy Hartse

Chair of the Supervisory Committee:
Martin Savage
Department of Physics

Quantum simulation provides a route to studying the real-time dynamics of strongly correlated quantum many-body systems beyond the reach of classical computation. This thesis focuses on two central obstacles that limit its practical utility: the high cost of implementing the required quantum circuits and the challenge of verifying their output in regimes where classical benchmarks are unavailable. Both issues are addressed by exploiting structure, either through prior physical knowledge of the target system or through special properties of selected quantum states, to reduce simulation overhead and to develop new techniques for quantum verification.

The background on which this work rests is assembled first: why classical simulation of quantum many-body systems is hard, why a quantum approach built on real-time evolution offers a way past these barriers, and the recurring tools the thesis relies on, including the thermalization of isolated quantum systems, the stabilizer formalism, and lattice gauge theories.

Next, known structure is exploited to reduce the cost of computing spectral densities and response functions, quantities at the heart of the description of nuclear reactions. Recasting the Gaussian integral transform of the response in a Fourier basis turns the moments that determine the spectrum into expectation values of real-time evolution, measurable directly on a quantum computer. Prior knowledge of

the low-order energy moments, the mean and variance, then allows the transform to be centered and rescaled so that far fewer evolution times are needed to resolve the response.

Turning from exploiting structure to discovering it, an exactly solvable form of weak ergodicity breaking is identified. In a $2+1$ -dimensional $\mathbb{Z}_2$ lattice gauge theory, a family of quantum many-body scars is shown to be composed entirely of stabilizer states. These scars carry exactly zero magic and exactly zero distillable entanglement, making them classically simple by both resource measures even as they sit within otherwise thermalizing, classically intractable dynamics, and their anomalous behavior is traced to an underlying commutant algebra.

This structure is then put to work as a verification tool. Because the stabilizer scars span a small, dynamically isolated subspace, their evolution can be computed exactly by classical means, and bounds on their non-stabilizerness ensure efficient direct fidelity estimation. Together these properties yield a self-contained and scalable benchmark: the measured fidelity of the scar states, compared against their exact classical prediction, quantifies the accuracy of the simulator and in turn bounds the fidelity of the genuinely intractable, non-scar dynamics run on the same device. The result is a concrete strategy for benchmarking quantum simulation at scale, in the regime where a classical check would otherwise be unavailable.

Finally, these results are drawn together, reflecting on what the shared strategy of exploiting structure achieves and where its limits lie, and sketching directions for extending it to richer systems and harder simulation tasks.

TABLE OF CONTENTS

	Page
List of Figures	iv
List of Tables	xiv
Chapter 1: Introduction	1
1.1 Quantum Simulation	1
1.2 Real-Time Dynamics	5
1.3 Dynamic Response Functions	6
1.4 Lattice Gauge Theories	7
1.5 Thermalization	10
1.6 Computational Complexity and Nonstabilizerness	16
1.7 Quantum Verification	19
Chapter 2: Faster Spectral Density Calculation Using Energy Moments . .	21
2.1 Introduction	21
2.2 Fourier Based Reconstruction	26
2.3 Gaussian Integral Transform	27
2.4 Numerical Examples	33
2.5 Discussion	41
2.A Derivation of Error Bounds	43
2.A.1 Periodic extension in the general case	44
2.A.2 Periodic extension with moment information	47
2.A.3 Truncation error	51
2.A.4 Statistical error	52
Chapter 3: Stabilizer Scars	55
3.1 Introduction	55

3.2	Model	56
3.3	Quantum Many-Body Scars	57
3.4	Entanglement and Magic Resource	59
3.5	Zero Distillable Entanglement in Lattice Gauge Theory Scar States	60
3.6	Generating Algebra	64
3.7	Experimental Realization.	65
3.8	Discussion	69
3.A	Lattice Gauge Theory and Ising Dual	70
3.B	Entanglement Structure and Symmetries	71
3.C	Derivation of the Scar Solutions	76
3.D	Details of the Quantum Many-body Scar Solutions	84
3.E	Other Lattice Geometries	93
3.F	Details of the Experimental Realization	93
3.G	Local Observables	97
3.H	Random Circuit Evolution	99
Chapter 4:	Benchmarking Quantum Simulation at Scale	102
4.1	Introduction	102
4.2	QMBS-Based Benchmarking	103
4.3	Efficient Direct Fidelity Estimation in a Stabilizer Scar Subspace	105
4.4	Fidelity Estimation for Classically Non-Simulable States	107
4.5	Prototype Model	107
4.6	Local Depolarization Errors are Scar-Subspace Agnostic.	108
4.7	Numerical Simulations	111
4.8	Discussion	116
4.A	Fidelity Estimation in Scar-subspace	117
4.A.1	Upper-Bound on α -SRE of ρ	117
4.A.2	Connection of α -SRE to DFE	119
4.B	QMBS Model Details	120
4.C	Average fidelity of the noisy channel	123
4.C.1	Bound tight near $p = 0$	124
4.C.2	Bound tight near $p = 1$	126
4.C.3	Concentration of Haar fidelity	133

4.D	MCMC Details	135
4.D.1	Metropolis-Hasting algorithm	135
4.D.2	I.i.d. MCMC samples	139
4.D.3	Simulating an experiment with shots	140
Chapter 5:	Conclusion and Outlook	145

LIST OF FIGURES

Figure Number		Page
1.1	<i>QMBS in the Rydberg Blockade and PXP Model.</i> (a) depicts the energy spectrum of an example non-integrable system. The dense region of the spectrum is comprised of the typical states, and the top and bottom sparse regions are the groundstate-like states. (b) is a plot from [60] and depicts the von Neumann entropy of the all the eigenstate of the PXP model. The x-axis is the energy of the eigenstate and the y-axis in the von Neumann entropy. The numbered states are the sub volume-law entanglement QMBS states and they correspond to the red lines in the energy spectrum in (a). The color corresponds to the density of states. (c) is a plot from [54] that depicts the oscillation of of a local observable as a function of time for the Néel state in the Rydberg blockade experiment. The y-axis is the domain-wall density, which corresponds to the density of number of neighboring atoms that are in the same state. This oscillation is also known as a revival. A non-scar would relax in about one coupling time (one Rabi or oscillation period), whereas the scar dynamics stay coherent for many such periods. This plot includes data from a 9 atom chain, a 51 atom chain, and a matrix product state (MPS) simulation.	14
2.1	<i>Response Functions.</i> The two model response functions shown as a function of frequency: black dots are for $S_A(\omega)$ while red squares for $S_B(\omega)$. The inset shows the same responses but on a narrower range and with linear scale. Also indicated are the mean μ and the square root of the variance σ for both distributions.	35
2.2	<i>Integral Transforms.</i> Comparison between the model response function $S_A(\omega)$ and $S_B(\omega)$ (black dots and red squares respectively) with their rescaled integral transforms $\delta\omega\Phi_A(\omega)$ and $\delta\omega\Phi_B(\omega)$ (green solid line and turquoise dotted line respectively). The inset shows the same curves but over the full energy range.	36

2.3	<i>Periodic Extension.</i> Comparison between the ideal (scaled) integral transforms $\delta\omega\Phi_B(\omega)$ (turquoise solid lines) and different periodic extensions. The bottom panel shows the approximation obtained using P for a generic response denoted as $\delta\omega\Phi_B^x(\omega)$ (blue dotted line) while the top panel shows the approximation obtained with a smaller P found using the known energy variance σ^2 and denoted as $[\sigma]\delta\omega\Phi_B^x(\omega)$ (brown solid line). The shaded grey areas are energies outside the interval $[\nu_{\min}, \nu_{\max}]$	38
2.4	<i>Error Scaling.</i> The top panel shows, as a function of the target error, the fractional reduction in the period from the general case value $P = P^{\text{gen}}$ obtained using information about the energy variance: the green line shows P_A^σ/P for model A while the turquoise is P_B^σ/P for model B . The bottom panel shows the observed approximation error for the two models as a function of the target error: the green line for model A and the turquoise line for model B . The solid black line represents the best error bound we derived for this approximation (see Section 2.A.2).	40
3.1	$\mathbb{Z}_2$ LGT. Degrees of freedom of $\mathbb{Z}_2$ LGT on the edges of a (periodic) two-dimensional square lattice; plaquette, electric field and Gauss law operators are shown.	57
3.2	<i>Quantum Many-Body Scars in $\mathbb{Z}_2$ LGT.</i> (a) Illustration of Dual Ising model with stabilizers $X_p X_{p+L}$ (eigenvalues in light and dark blue) and $Z_p Z_{p+L}$ (eigenvalues in light and dark red). (b) QMBS scar solutions of a $L \times 2$ lattice where L is even, shown in terms of $X_p X_{p+L}$ and $Z_p Z_{p+L}$ eigenvalues. (c) The scar subspace for odd L is spanned by the states shown.	58
3.3	<i>Quantum Many-Body Scars: Lattice Gauge Theory vs. Ising Dual.</i> Top row: Distillable entanglement of all eigenstates for even L , for $L = 8$ and $g = 0.9$. The scar states, given by Eq. (3.3) and shown as red stars, appear in both the LGT (left panel) and the dual Ising model (right panel) and are characterized by exactly zero distillable entanglement. Bottom row: For odd L ($L = 7$ and $g = 0.9$), the scar eigenstates form linear superpositions of the stabilizer states, Eq. (3.4). The scars exhibit zero distillable entanglement in the LGT, but not in the dual Ising model.	61

3.4 *Quantum Many-Body Scars: Lattice Gauge Theory vs. Ising Dual II.* (a) Superselection sectors, $\sigma_\nu^z = \pm 1$ on each boundary link $\nu \in \{\nu_1^a, \nu_2^a, \nu_1^b, \nu_2^b\}$, are associated with Gauss laws. (b) An additional superselection sector, $\prod_{\nu \sim \mathcal{L}_x^A} \sigma_\nu^z \sim Z_a Z_b$, where $\mathcal{L}_x^A = \mathcal{L}_x \cap A$, arises from restricting $\prod_{\nu \sim \mathcal{L}_x} \sigma_\nu^z = V_x$ to A (similar for $\bar{A}$). (c) The $Z_p Z_{p+L}$ eigenvalue configuration, for $p \leq a$ and $p \geq b$, is shown for the state $|\psi_{+1,-1}\rangle$ from Eq. (3.11), with a defect trapped inside A . Along with the $X_p X_{p+L}$ and Gauss law constraints, the reduced density matrix of $|\psi_{+1,-1}\rangle$ on $\bar{A}$ (and A) has at most rank 2. The light red and dark red colors indicate the $Z_p Z_{p+L}$ eigenvalue as in fig 3.2. (d) While $Z_p Z_{p+L}$ constrains horizontal links in the LGT, Gauss law eigenstates require vertically separated vertical links to be in either a $+1, -1$ or $-1, +1$ configuration, providing an effective spin-1/2 degree of freedom. Imposing the $X_p X_{p+L}$ constraint leads to two states—one even, one odd under $\prod_{\nu \sim \mathcal{L}_x^A} \sigma_\nu^z$ —resembling the degenerate ground state of a 1D Ising model without a transverse field. 63

3.5 *QMBS Preparation.* (a) Example of a scar basis state specified by the eigenvalues of $X_p X_{p+L}$ and $Z_p Z_{p+L}$. (b) Initial state for the preparation algorithm in the LGT formulation ($V_y = 1$). Horizontal links are initialized in σ^z (electric) eigenstates, while vertical links are initially set to $|0\rangle$. (c) The state preparation circuit operates on pairs of vertically separated qubits located at ν and $\nu + \hat{y}$. An example is shown where the circuit prepares a state with $\sigma_\nu^x \sigma_{\nu+\hat{y}}^x = -1$ and $\sigma_\nu^z \sigma_{\nu+\hat{y}}^z = -1$; combined with the horizontal links, this yields a Gauss law $+1$ eigenstate. Minor circuit adjustments enable preparation of any combination of $\sigma_\nu^x \sigma_{\nu+\hat{y}}^x = \pm 1$ and $\sigma_\nu^z \sigma_{\nu+\hat{y}}^z = \pm 1$ states, allowing access to any scar subspace basis state. The circuit maintains product-state structure along L ; each state remains a 2-qubit product state. 67

3.6	<i>Experimental Realization.</i> (a) The absolute value of the Loschmidt echo, $ L(t) $, is shown as a function of time, comparing an initial state within the scar subspace (red lines) to a thermalizing state (blue lines) for system size $L = 11$. The time-averaged value at late times is indicated by the dark red dashed line for the scar state. The scar echo decays as $ L(t) \sim 1/\sqrt{4L}$ (gray dotted line), where $4L$ is the size of the scar subspace. In contrast, the thermalizing state (blue lines) explores a much larger Hilbert space. For a random circuit that fully explores this space, $ L(t) $ would decay to $ L(t) \sim 1/\sqrt{2^{2L-1}}$ (gray dotted lines), however, due to energy conservation, the thermalizing state's late-time average (dark blue dashed line) remains slightly above this value. (b) The distillable entanglement $S_{\text{vN}}^{\text{dist}}$ of a time-evolved scar state (red) remains zero throughout. By contrast, for a thermalizing state (light blue), $S_{\text{vN}}^{\text{dist}}$ increases over time until reaching a constant, non-zero saturation value (blue dashed line).	68
3.7	<i>Symmetries of the LGT.</i> (a) Shown are $2L$ Gauss laws, Eq. (3.13), at sites s , which are local symmetry operators that commute with the Hamiltonian H , thereby defining superselection sectors. (b) For periodic boundary conditions, the ribbon operators V_x and V_y , Eq. (3.14), defined along $\mathcal{L}_x$ and $\mathcal{L}_y$, respectively, specify additional superselection sectors.	73
3.8	<i>Analytic Derivation.</i> (a) The derivation of the scar eigenstates presented in this work begins with the degrees of freedom in the $\mathbb{Z}_2$ lattice gauge theory (LGT), represented by spins located on the links of a $L \times 2$ -site lattice with periodic boundary conditions. (b) This system is then mapped onto a dual Ising model, subject to the parity constraint $\prod_p X_p = 1$, where spin-1/2 degrees of freedom reside at the centers of the plaquettes (denoted by $p = 1, \dots, 2L$). A Jordan-Wigner transformation (indicated by the green arrow) initiates in the top row (covering plaquettes $1, \dots, L$) and extends to the bottom row (from sites $L + 1$ to $2L$), transforming the system into a two-component Majorana representation. (c) The Majorana modes are combined vertically to form a two-flavor fermion representation, where $p = 1, \dots, L$ denotes one-dimensional sites. This representation enables the identification of scar subspaces for both even and odd L via the commutant of the algebra generated by the terms in H .	77

- 3.9 *Distillable Entanglement for Varying Couplings.* The following plots show $S_{\text{vN}}^{\text{dist}}$ for a range of couplings. Panels in the top row display results for an even system size $L = 8$, while the bottom row shows results for an odd system size $L = 7$. The left panels correspond to results from the LGT formulation, and the right panels to those from the Ising model. Scar states are highlighted with red stars. In the Ising dual formulation, scar states exhibit non-zero distillable entanglement, blending with other typical states and becoming indistinguishable. However, in the LGT, scar states maintain exactly zero distillable entanglement, which differentiates them clearly. While the limits $g \rightarrow 0$ and $g \rightarrow \infty$ yield ground states that are straightforward to understand, the behavior of mid-spectrum states is complex, with scar eigenstates persisting in these limits. 86
- 3.10 *Contributions to von Neumann Entanglement Entropy.* The top row shows the distillable entanglement $S_{\text{vN}}^{\text{dist}}$, the middle row the symmetry entanglement $S_{\text{vN}}^{\text{symm}} \equiv -\sum_s p_s \log(p_s)$, and the bottom row the total von Neumann entropy $S_{\text{vN}} = S_{\text{vN}}^{\text{dist}} + S_{\text{vN}}^{\text{symm}}$. The left column contains results from the LGT, while the right column shows results for the Ising dual. Quantum many-body scar (QMBS) states are highlighted in red, with the maximum possible symmetry entanglement indicated by horizontal gray dashed lines. For $g = 0.9$, these results show that QMBS states can be clearly distinguished by their distillable entanglement in the LGT, where they exhibit exactly zero values. By contrast, the total von Neumann entropy, although widely used, is finite and therefore less reliable for identifying QMBS states. Indeed, at other coupling values and in smaller systems, scar states may blend with typical states in terms of S_{vN} even within the LGT, rendering them indistinguishable. 90
- 3.11 *Total von Neumann Entanglement Entropy.* Shown is the (total) Neumann Entropy S_{vN} for lattice size $L = 7$ with different couplings and with the scars, identified via their vanishing *distillable* entanglement, highlighted in red. (a) $g = 0.1$. (b) $g = 0.9$, (c) $g = 1.5$, (d) $g = 2.5$. Especially at high and low coupling, panels (a) and (d), the spectrum starts separating into energy bands and identifying the scars becomes impossible based on S_{vN} when not separated into distillable and symmetry components. For smaller sized systems, distinguishing scars from non-scars based on S_{vN} , instead of via the distillable entanglement, becomes impossible even at intermediate couplings. 92

3.12 *Distillable Entanglement of Eigenstates for Lattice Geometries $L \times k$, where $k > 2$.* ((a) Left Panel): No low entanglement states are observed for 4×4 (here shown for $g = 0.9$), and for ((b) Right Panel): 5×3 . While our analytical construction clearly does not extend to these geometries, this numerical study does not strictly rule out the possibility of different types of QMBS in these geometries, for instance potentially appearing only near specific coupling values. 94

3.13 *Alternative Circuit for QMBS Preparation Using Mid-Circuit Measurements.* (a) Initial state for the preparation algorithm in the LGT formulation ($V_y = 1$). Horizontal links are in σ^z (electric) eigenstates, and vertical links in σ^x (magnetic) eigenstates. Alternating $|0\rangle$ and $|1\rangle$ (or $|+\rangle$ and $|-\rangle$) yields eigenstates of any combination of stabilizers $Z_p Z_{p+L}$ and $X_p X_{p+L}$. (b) Mid-circuit measurement-based scheme for simultaneous Gauss law eigenstate preparation using L ancillas. (c) The circuit measures the combined σ^z parity of two neighboring vertical links via controlled- σ^z ($\Gamma(\sigma^z)$) operations and an ancilla a . The outcome dictates controlled-NOT operations on a vertical qubit, correcting Gauss laws in upper and lower rows, simultaneously. Control instructions classically depend on the desired target scar (via the horizontal links' state) and V_y sector. (d) A different starting point is shown for the scar state-preparation circuit, a simultaneous eigenstate of the stabilizer $Z_p Z_{p+L}$ and the Gauss law constraints. (e) This initial states is not eigenstates of an $X_p X_{p+L}$; to transform it into asimultaneous eigenstate, a circuit is used involving mid-circuit measurements and L ancillas. In a single shot, it measures $\sigma_\nu^x \sigma_{\nu'}^x$ values, where ν and ν' denote vertically stacked, vertically oriented links. From this the value of all $X_p X_{p+L}$ is determined. Based on the measurement outcome and the desired target state, a controlled σ^z operation is applied. 95

3.14 *Loschmidt Echo in the Scar Subspace for Large System Sizes L .* The plot shows the absolute value of the Loschmidt echo, $|L(t)|$, evaluated for scar initial states with system sizes up to $L = 901$. The echo exhibits a power-law decay, $\sim (gt)^{-1/2}$, approaching an asymptotic value of $1/\sqrt{4L}$ at long times. The fit error is determined by varying the fit range and L 97

3.15 *Local Observables.* The expectation values of electric operators, $\langle \sigma_\nu^z \rangle$, for links oriented horizontally (top panel) and vertically (bottom panel), contrasting scar initial states (red and orange lines) with thermalizing non-scar initial states (dark and light blue lines) for $L = 9$ and $g = 0.9$. For non-scar initial states, the expectation values thermalize to a constant, asymptotic value (indicated by dashed blue lines). In contrast, the expectation values within the scar subspace exhibit oscillatory behavior, albeit without perfect revivals, a consequence of the non-equidistant energy spacing of scar eigenstates. This limits the utility of local observables, as distinguishing these oscillations from finite-volume effects and experimental imperfections may prove challenging. We argue that observables such as the Loschmidt echo or the distillable entanglement, however, allow a stronger distinction between scar and non-scar dynamics. 98

3.16 *Experimental Realization with a Random Circuit.* (a) Evolution of scar and non-scar states through random circuits, as defined in Eq. (3.55). The absolute value of the Loschmidt echo, $|L(s)|$, is shown as a function of the number of circuit steps s , contrasting an initial state within the scar subspace (red lines) with a state outside this subspace (blue lines) for a system size $L = 11$. Dashed lines represents late-time averages. For the scar state, the echo decays to $|L(s)| \sim 1/\sqrt{4L}$ (gray dotted line), where $4L$ is the dimension of the scar subspace. Conversely, the non-scar state decays to $|L(s)| \sim 1/\sqrt{2^{2L-1}}$ (gray dotted line), reflecting the lack of an energy conservation constraint. (b) The distillable entanglement, $S_{\text{vN}}^{\text{dist}}$, for scar versus non-scar initial states exhibits similar behavior to that shown in the main text for Hamiltonian evolution. 100

- 4.1 *Benchmarking Quantum Simulations at Scale.* A protocol for benchmarking the fidelity of quantum simulations of non-equilibrium dynamics in regimes beyond classical simulability. We consider models which are non-integrable, for which typical states, $|\psi_{\mathcal{H}}\rangle$, thermalize. Such dynamics are not efficiently classically simulable, and benchmarking performance for such states is intractable. For models that host stabilizer scars, under consideration here, we can exploit the fact that for initial states in this subspace, $|\psi_{\mathcal{S}}\rangle$ dynamics are classically simulable and we can perform direct fidelity estimation efficiently. Assuming local depolarizing noise, we demonstrate for a specific stabilizer scar model [149] that the fidelity of classically non-simulable states is approximately equal to that of the states restricted to the scar subspace, thus providing a method to verify performance for arbitrary inputs. 104
- 4.2 *Noise Acting on Scar vs. Non-Scar States.* Comparison between the fidelity under a local depolarizing channel, averaged over Haar-random input from the full Hilbert space, $F(\mu_{\mathcal{H}})$ (solid lines), and bounds of the respective fidelity for scar states $F(\mu_{\mathcal{S}})$ (shaded bands, Eq. (4.15) and App 4.C). The difference, shown in the lower panel, vanishes as $n \rightarrow \infty$ or for small p , and (trivially) for $p \rightarrow 1$. The tighter bound is shown, and the crossover between where Eq. (4.15) is tighter and where the App result applies is marked by square markers. 109
- 4.3 *Numerical MCMC Demonstration.* (a) Failure probability to reach precision $\epsilon = 5 \cdot 10^{-2}$ of MCMC-based DFE between a target $\rho \in \mathcal{S}$ and a reference state $\sigma \in \mathcal{S}$ as a function of the number of sampled Pauli operators with exact measurements. Results shown for $L = 11, 31, 51, 71$ ($n = 22, 62, 102, 142$ qubits) compared with the bound $s = \lceil 1/(\epsilon^2 \delta) \rceil$ [110] are independent of system size when expectation values are known exactly. (b) Finite-shot estimation. For $\rho \in \mathcal{S}$, we show the maximal shot number $\langle N_{\rho} \rangle$ (dashed black) versus the optimal shot cost (points; solid black fit), attainable if exact expectation values were known (see App 4.D) for precision $\epsilon = 5 \cdot 10^{-2}$ and failure probability $\delta = 1 \cdot 10^{-2}$. Shown are the average shot number over $s = 10^4$ Paulis, $\langle N_{\rho} \rangle = N_{\rho}/10^4$ samples. Both bounds scale as $\sim d_{\mathcal{S}}^2$ 112

- 4.4 *Circuit Simulation.* Quantum circuit used to simulate Eq. (4.11), consisting of state preparation, (random or Trotterized) time evolution, and DFE. Local depolarizing noise is assumed, except during state preparation and measurement. The state preparation circuit consists of $L = n/2$ two-qubit Bell-pair sub-circuits and is applied to an input computational basis product state $|s\rangle$, where s is an n -qubit bitstring. For special s , the circuit prepares QMBS basis states, while for all others it produces orthogonal non-scar states. 114
- 4.5 *Circuit Simulation.* (a) Classical emulation of the circuit in Fig. 4.4 for $n = 10$ qubits comparing fidelities as a function of p for different circuit depths ℓ , comparing scar basis state (dashed lines) with non-scar inputs (dotted lines). The results are compared with the analytical expression for the *average* fidelity, Eq. (4.9), modified by an effective error rate $p_{\text{eff}} = 1 - (1 - p)^\ell$ accounting for the ℓ layers (solid black lines). The bottom panel shows the difference $\Delta F \equiv F_S - F_H$, decreasing with ℓ . (b) Noisy circuit simulation with $n = 10$ qubits, $\ell = 15$ layers and $p = 0.001$, where Pauli expectation values are estimated from finite-shot sampling. Top: Fidelity versus the number of MCMC samples, s . Circles show shot-sampled results (10^4 shots per Pauli observable), while solid lines correspond to exact Pauli expectation values. Two representative MCMC chains are shown out of 100. Bottom: Final fidelity precision at the end of shot-sampled MCMC chains ($s = 10^4$). 115
- 4.6 *Stabilizer scar subspace of the prototype model.* (a) The model, Eq. (4.11) of the main text, is defined on an $L \times 2$ lattice with periodic boundary conditions. The dual lattice of the LGT, Eq. (4.29), is shown in gray, in which the degrees of freedom reside on the links. (b) For odd L , the QMBS subspace is spanned by the orthonormal stabilizer basis states shown here, each characterized by a distinct eigenvalue pattern of the stabilizer generators $X_i X_{i+L}$ (blue) and $Z_i Z_{i+L}$ (red), where darker (lighter) shading indicates a +1 (−1) eigenvalue. (c) Initial states are prepared using simple Bell circuits specified by an input bit string s , which can generate any QMBS basis state or a state outside the scar subspace with the same circuit complexity. Subfigures (a) and (b) adapted from Ref. [149]. 121

- 4.7 *MCMC suggest-step decision tree.* Shown is the decision tree used to uniformly generate Pauli operators with nonzero expectation value within the stabilizer scar subspace. This proceeds in four steps: (i) a Z operator pattern is generated, in which either a “defect hops” (left column) or does not (right column); (ii) this Z pattern is further dressed by a random product of $\{\mathbb{I}_i \mathbb{I}_{i+L}, Z_i Z_{i+L}\}^{\otimes L}$ operators; (iii) based on the resulting Z pattern, an X pattern is generated that either connects the defect’s origin and target (left column, not shown are the inverted X patterns in the top row) or merely changes its flavor (right column); and (iv) this X pattern is similarly dressed by a random product of $\{\mathbb{I}_i \mathbb{I}_{i+L}, X_i X_{i+L}\}^{\otimes L}$ operators. The corresponding matrix elements, together with the phase obtained from Eq. (4.77), are tracked throughout to construct a $4L \times 4L$ matrix in the QMBS subspace. 137
- 4.8 *Additional noisy circuit simulation with shots.* (a) Results for $n = 10$ qubits, $\ell = 15$ circuit layers, and depolarizing noise strength $p = 0.001$, analogous to Fig. 4.5(b) of the main text. Eight different randomly selected MCMC chains are shown out of 100 generated. Circles denote shot-sampled estimates (10^4 shots per Pauli observable), while solid lines show the corresponding exact Pauli expectation values. Some shot-sampled trajectories exhibit an apparent systematic offset from the exact curves. This arises from the construction of the statistical uncertainty: the error bars represent the accumulated uncertainty of the fidelity estimate after s MCMC samples, obtained via error propagation. As a result, an early (or intermediate) statistical outlier in a Pauli expectation value can bias the cumulative estimate until its effect is gradually averaged out by subsequent samples. (b) Same as the bottom panel of Fig. 4.5(b) of the main text, showing the final fidelity precision after $s = 10^4$ MCMC samples. 142
- 4.9 *Shot dependence.* The same MCMC chain as in Fig. 4.5(b) of the main text applied to noisy circuit simulations with $n = 10$ qubits, $\ell = 15$ layers, and depolarizing noise strength $p = 10^{-3}$. The number of shots per Pauli observable is varied between 10^3 (10% of the shots used in the main text, left panel), $5 \cdot 10^3$ (50%, middle panel), and $5 \cdot 10^4$ (500%, right panel). The bottom panels show the resulting precision after $s = 10^4$ MCMC samples. 143

LIST OF TABLES

Table Number		Page
3.1	<i>Mapping Between the $\mathbb{Z}_2$ LGT and its Ising dual.</i> The degrees of freedom are spin- $\frac{1}{2}$ variables, which reside on the links ν in the LGT and on the faces of the plaquettes p in the dual formulation. The correspondence between magnetic, electric, and Gauss law terms is outlined. In the dual picture, the Gauss laws—originally local constraints in the LGT—is eliminated. The parity constraint in the dual corresponds to the global identity that the product of all plaquettes equals one in the LGT. Additionally, the mapping of the $Z_p Z_{p+L}$ and $X_p X_{p+L}$ stabilizers is presented.	72

ACKNOWLEDGMENTS

This thesis would not have been possible without the support and guidance of many people, and I will do my best to thank them all here. I want to start by thanking my advisor, Professor Martin Savage. He gave me the freedom to explore what I found interesting, and gave me many opportunities to learn and grow as a researcher.

I have learned so much from everyone in the Incubator for Quantum Simulation (IQuS) group and I want to say thank you to all of the current and former members. When I was in my first year of graduate school, Professor Alessandro Roggero advised and worked with me on my first project and taught me the research process. I also had the privilege of working closely with Professor Niklas Mueller over the past three years. He gave me invaluable guidance and helped me grow confidence as a researcher.

I would like to thank Zhiyao Li, Ivan Chernyshev, Nikita Zemlevskiy, and Henry Froland, the four other graduate students who joined IQuS alongside me, for sharing this journey and making it far more enjoyable. I also wanted to thank Katie Hennessey for her invaluable work as an administrator for our group, where she organized workshops and made travel go smoothly.

I would like to thank my reading committee members, Professors Lukasz Fidkowski and Silas Beane, who took the time to advise me on this thesis. I also would like to thank, my other supervisory committee members, Professors Stephen Sharpe, Subhadeep Gupta, Rahul Trivedi and Paul Beame, who took the time to serve on my committee.

I would also like to thank some close friends that have supported me during my

time in graduate school: Jordan, Nic, Lou, Matt, Rachel, CJ, Niall, Alex, Sam, Sarah, and many others. I made many friends at the Edgeworks climbing gym in Ballard, which was a great place to blow off steam and build a community. Washington has been an amazing place to live and I am very thankful to have had the opportunity to spend time hiking, skiing, climbing, and running around this beautiful state.

I would like to thank my family for their support over the years, including my parents, Hans and Peggy, my sister, Sara, and my grandmother, Lois. Finally, I want to thank my fiancée, Adina. She has always been there to support me and give me advice when I need it, and I can't imagine I would have made it far without her.

DEDICATION

To my parents, Hans and Peggy, and my fiancée, Adina.

Chapter 1

INTRODUCTION

1.1 *Quantum Simulation*

The starting point for understanding why simulating quantum many-body systems is hard is the exponential growth of Hilbert space. A system of n two-level constituents occupies a Hilbert space of dimension 2^n . A generic state $|\psi\rangle = \sum_x c_x |x\rangle$ is specified by 2^n complex amplitudes. The time evolution of this state is the action of a $2^n \times 2^n$ unitary. Writing down the state, let alone evolving it, becomes intractable quickly. However, the exponential size of Hilbert space is not by itself the obstacle, since there are many quantum states that are easy to simulate. The difficulty arises only when a state uses this space generically, in a way that resists any compact classical description. Whether a state does so is governed by more than one property, its entanglement, but also its nonstabilizerness, or magic, and a central theme of this chapter is that these are distinct axes: a state can be highly entangled yet classically simple. Which states carry such exploitable structure, and which do not, is the question that separates tractable quantum systems from intractable ones.

Tensor-network methods such as the density-matrix renormalization group represent the limited-entanglement ground states of local Hamiltonians with only polynomial resources [1], and quantum Monte Carlo evaluates equilibrium properties by stochastic sampling [2].

Lattice field theory, and lattice quantum chromodynamics in particular, is a canonical success of Markov chain Monte Carlo. Discretizing spacetime and working in imaginary (Euclidean) time gives the path integral a real, positive Boltzmann weight e^{-S_E} that can be treated as a probability and sampled [3]. This approach has enabled

first-principles calculations of the hadron spectrum, decay constants, form factors, and the equation of state of QCD at zero baryon density [4]. The method fails whenever this probabilistic interpretation breaks down. When the Boltzmann weight stops being positive, calculations incur the Monte Carlo *sign problem*, in which the sampled quantity oscillates in sign and the statistical error grows exponentially. At finite baryon chemical potential the fermion determinant becomes complex, destroying the positive probability measure required for importance sampling, limiting reliable lattice calculations to the regime of small chemical potentials. As a result, the dense region of the QCD phase diagram, including the conjectured critical point and the matter inside neutron stars, remains largely inaccessible [5]. Real-time evolution amplifies this problem by rotating back to Minkowski time, thus turning the entire Boltzmann weight into a purely oscillating phase e^{iS} , meaning transport coefficients, spectral functions, and out-of-equilibrium dynamics cannot be computed directly [6]. The sign problem is NP-hard in general, so no universal solution is expected [7]. Tensor-network methods face a different barrier. Although they provide efficient descriptions of many one-dimensional systems, their computational cost increases rapidly in higher spatial dimensions, and under generic non-equilibrium evolution the entanglement of $|\psi(t)\rangle$ grows linearly in time, forcing the bond dimension to grow exponentially. Classical simulations are therefore restricted to relatively short times. [1]. Consequently, many of the most physically compelling regimes, such as real-time dynamics, scattering, thermalization, transport, and finite-density matter, are precisely those in which classical methods break down.

Famously, Feynman's resolution to this impasse was to simulate a quantum system using another, controllable quantum system. In this way, nature itself, rather than a classical processor, could bear the exponential complexity [8, 9]. Lloyd made the idea concrete by showing that the dynamics of any local Hamiltonian can be reproduced efficiently on a quantum computer [10], and the intervening decades have turned the proposal into a programmable reality across platforms ranging from trapped ions and

superconducting qubits to neutral-atom and Rydberg arrays. At the heart of these platforms is their ability to implement real-time evolution,

$$|\psi(t)\rangle = e^{-iHt} |\psi(0)\rangle, \quad (1.1)$$

the very operation that classical methods cannot sustain.

Nuclear and high-energy physics lie squarely in the regime where this capability is most needed [11, 12, 13]. These fields are governed by strongly-coupled quantum field theories whose nonperturbative content is, at equilibrium, accessible through Euclidean lattice Monte Carlo, but whose real-time, finite-density, and out-of-equilibrium physics is obstructed by the sign problem. Problems such as the thermalization of matter produced in heavy-ion collisions, transport and parton dynamics, and inclusive scattering cross sections lie beyond its reach. Lattice gauge theories provide the natural framework and a prime target for quantum simulation [14], with even Abelian toy models such as the $\mathbb{Z}_2$ and $U(1)$ theories capturing confinement, string breaking, and thermalization while serving as tractable testbeds. In low-energy nuclear physics, a paradigmatic class of observables is the set of response functions that govern inclusive scattering, for instance in neutrino–nucleus and electron–nucleus reactions. Response functions encode the dynamical linear response of the nucleus and, reduce to the spectral density of the nuclear Hamiltonian, a quantity naturally extracted from real-time evolution, Eq. (1.1) [15, 16]. Despite these diverse applications, the device is asked to do the same thing: prepare a state, evolve it in real time, and infer physics from the evolved state.

The decades since Lloyd’s proposal have seen quantum simulation move from principle to practice, with programmable devices now operating across several architectures, each carrying its own tradeoffs for simulating many-body physics. Superconducting processors offer fast gates and a mature fabrication path, but their qubits have relatively short coherence times and, more importantly, are wired with fixed low-degree connectivity, so the nonlocal interactions typical of gauge and nuclear models

must be synthesized through costly networks of SWAP gates [17, 18]. Trapped-ion systems invert this balance: they provide long coherence, some of the highest gate fidelities demonstrated, and native all-to-all connectivity well suited to nonlocal couplings, at the price of slower gates and the difficulty of scaling a single register to large sizes [19, 20]. Neutral-atom and Rydberg arrays reach the largest system sizes and allow the qubit geometry itself to be reconfigured, though their two-qubit gates are slower and historically lower in fidelity [21, 22, 23]. The steady improvement in qubit number, gate fidelity, and coherence is what has made the physics targets of this thesis experimentally relevant rather than purely theoretical.

That quantum hardware can, in principle, simulate these systems does not by itself deliver usable physics, because the devices that exist operate under real constraints. Present-day processors are noisy and of limited scale, the noisy intermediate-scale quantum (NISQ) regime, in which circuit depth is bounded by accumulating gate errors. Quantum error correction offers the eventual route beyond this limit: a logical qubit is encoded redundantly across many physical qubits, and the threshold theorem guarantees that once physical error rates fall below a fixed threshold, logical errors can be suppressed arbitrarily by enlarging the code [24, 25]. Yet even ideal logical qubits would not make these simulations free. Real-time evolution remains intrinsically expensive, and its output remains hard to trust, and it is these two obstacles, cost and verification, that this thesis sets out to address.

By cost we mean the total physical resources a simulation consumes to reach a target accuracy: the number of qubits, the circuit depth and total gate count, the number of measurement repetitions, and the total duration of real-time evolution. The existence of an efficient algorithm is only an asymptotic statement about how these scale, and the concrete resources demanded at the system sizes and precisions of physical interest can be prohibitive, especially on the resource-limited devices of the near-term and early fault-tolerant eras. Fundamental limits make some of this cost unavoidable. The no-fast-forwarding theorem establishes that a generic Hamiltonian

cannot be simulated to time T with cost growing more slowly than T [26], so the long-time dynamics of a typical system are intrinsically expensive. This limit is generic rather than universal: Hamiltonians with special structure, such as commuting or free-fermion systems, can be fast-forwarded and evolved to late times at greatly reduced cost. Where such structure is absent, however, long-time evolution remains costly, and making quantum simulation practical requires algorithms that remove every avoidable resource from the calculation.

The second major challenge is verification. Suppose a quantum processor returns an answer for a quantity that no classical computer can reach. How is its correctness to be established? The very classical intractability that justifies running the problem on a quantum computer is what makes its output hard to check, and on real hardware noise can quietly corrupt the result. In fact noisy quantum circuits without error correction are, in important regimes, efficiently classically simulable, as shown with noisy random circuit sampling at any constant noise rate by Aharonov et al [27]. Efficient classical algorithms already capture many observables, so a genuine demonstration must lie beyond them. Verifying a classically infeasible computation can be as costly as performing it and certification must be robust to noise [28, 29, 30, 31, 32]. Demonstrating that a quantum simulation is simultaneously correct, robust to noise, and beyond classical reach remains one of the central challenges facing the field.

1.2 Real-Time Dynamics

The sign problem introduced in Sec. 1.1 is most severe for real-time dynamics. Evaluating the propagator e^{-iHt} by stochastic sampling fails because the path-integral weight oscillates rather than decays, destroying the probabilistic interpretation required for Monte Carlo methods. This is the setting in which a quantum computer is most naturally suited, since a controllable quantum system can enact the unitary e^{-iHt} of another directly rather than sampling it.

For a local Hamiltonian written as a sum of terms, $H = \sum_j H_j$, the evolution can be approximated by a product formula that interleaves the easily implementable factors $e^{-iH_j\delta t}$, the Trotter–Suzuki decomposition [10, 33]. The first-order Trotter formula is $e^{-iHt} \approx \left(\prod_j e^{-iH_j t/n}\right)^n$, whose error arises from the noncommutativity of the H_j terms and scales as $\mathcal{O}(t^2/n)$. Improving the accuracy therefore requires increasing the number of Trotter steps and, consequently, the circuit depth. More advanced approaches based on linear combinations of unitaries and qubitization offer improved asymptotic scaling [34, 35]. Regardless of the simulation algorithm, the computational cost grows with the evolution time, the desired precision, and the system size. Measured in gate count or queries to the time-evolution operator, long-time dynamics are therefore intrinsically expensive. Controlling that cost is one of the two central concerns of this thesis.

Real-time evolution is not merely the mechanism by which a quantum simulator operates; many physical quantities of interest are themselves expectation values of e^{-iHt} . The dynamic response functions introduced in Sec. 1.3 are Fourier transforms of real-time correlation functions, making their reconstruction a problem of estimating matrix elements of the evolution operator. Additionally, the Loschmidt echo $|\langle\psi|e^{-iHt}|\psi\rangle|^2$ that diagnoses the scars of Chapter 3 is a real-time quantity.

1.3 Dynamic Response Functions

Many of the most important observables in nuclear and high-energy physics are not static ground-state properties but dynamic response functions, which characterize how a system responds to an external probe such as a scattered electron or neutrino [36]. A response function takes the form

$$S(\omega) = \sum_m |\langle\psi_0|\hat{O}|\psi_m\rangle|^2 \delta(\omega - (E_m - E_0)), \quad (1.2)$$

a sum over the excited states $|\psi_m\rangle$ reached from the ground state $|\psi_0\rangle$ by an excitation operator $\hat{O}$, resolved according to the energy transfer ω . The electromagnetic and

weak response of nuclei, which governs processes from electron scattering to neutrino detection, is of exactly this form.

Computing $S(\omega)$ directly is generally intractable. It requires resolving a large number of excited states across what is often a continuum, and the underlying real-time correlator reinstates the sign problem for classical Monte Carlo methods. A standard way around the full spectrum is to compute an integral transform of the response, in which the sharp delta functions of Eq. (1.2) are smeared against a kernel to give a smooth function that is easier to evaluate and can afterward be inverted [37, 15, 38]. Among the most widely used are the Lorentz and Gaussian integral transforms.

On a quantum computer these integral transforms arise naturally because the smeared response can be written directly in terms of real-time evolution. In the Gaussian integral transform the kernel is a Gaussian, and expanding it in a Fourier basis recasts the transformed response as a finite set of expectation values of e^{-iHt} evaluated at discrete times [15, 38]. The resolution of the reconstructed spectrum improves as more moments, and hence more evolution times, are included. Here the central challenge becomes computational cost. Improving the spectral resolution requires estimating an increasing number of Fourier moments, each demanding an independent real-time evolution. As the target resolution sharpens, the required evolution times and measurement overhead grow rapidly. Chapter 2 develops a method for reducing exactly this measurement overhead by incorporating prior information about the spectral density into the reconstruction.

1.4 *Lattice Gauge Theories*

Lattice gauge theories are among the principal targets of quantum simulation because they provide a nonperturbative formulation of strongly coupled quantum field theories such as QCD. Whereas classical lattice calculations typically employ the Euclidean path-integral formulation, quantum simulators naturally work with a Hamiltonian formulation, the first of which was introduced by Kogut and Susskind [39, 40]. In this

formulation the gauge degrees of freedom reside on the links of a spatial lattice and are valued in a gauge group G , and matter, where present, resides on the vertices. At each vertex a local gauge transformation G_v generates the gauge symmetry, and the physical states are those left invariant by it, obeying Gauss's law,

$$G_v |\psi\rangle = |\psi\rangle \quad \text{for every vertex } v. \quad (1.3)$$

The physical Hilbert space is therefore the gauge-invariant subspace rather than the full tensor-product space of link variables. This constraint is the organizing feature of the gauge-theory chapters of this thesis: it shapes the spectrum and underlies the scars of Chapter 3, and, as discussed below, it dictates how entanglement is defined.

A simple example, and the setting of this thesis, is $\mathbb{Z}_2$ gauge theory in two dimensions with a single qubit on each link, e , representing the gauge field. In the standard convention the Gauss-law operator at a vertex is the star $G_v = \prod_{e \ni v} \sigma_e^z$, the product of σ^z over the links meeting at v , while the magnetic flux through a plaquette is measured by $B_p = \prod_{e \in p} \sigma_e^x$. The physical states satisfy $G_v |\psi\rangle = |\psi\rangle$. The dynamics are governed by a Hamiltonian containing both magnetic-flux and electric-field terms, whose relative strength is controlled by the coupling g :

$$H_{\mathbb{Z}_2} = - \sum_p B_p - g \sum_e \sigma_e^z \quad (1.4)$$

At $g \rightarrow 0$, the ground state is the toric code ground state [41], a topologically ordered stabilizer state, and varying g drives a confinement–deconfinement transition.

Beyond constraining the dynamics, Gauss's law also means that the physical Hilbert space no longer factorizes across spatial regions. The non-factorization implied by Eq. (1.3) makes entanglement in a gauge theory subtler than in an ordinary system. In a conventional lattice model the Hilbert space factorizes across a spatial cut, $\mathcal{H} = \mathcal{H}_A \otimes \mathcal{H}_B$, and the entanglement entropy $S_A = -\text{Tr } \rho_A \ln \rho_A$ of the reduced state $\rho_A = \text{Tr}_B |\psi\rangle\langle\psi|$ is unambiguous. A gauge theory does not admit such a factorization. Gauss's law constrains the links crossing the boundary and ties the two

regions together, so that $\mathcal{H}_{\text{phys}} \neq \mathcal{H}_A \otimes \mathcal{H}_B$ [42, 43, 44, 45]. The bipartite entropy then depends on how $\mathcal{H}_{\text{phys}}$ is embedded into a tensor-product space, equivalently on a choice of boundary operator algebra, and ceases to be a single operationally defined number.

Choosing the algebra of gauge-invariant operators resolves the ambiguity and yields a canonical decomposition of the entropy [43, 44, 45],

$$S_A = - \sum_{\vec{r}} p_{\vec{r}} \ln \frac{p_{\vec{r}}}{d_{\vec{r}}} + \sum_{\vec{r}} p_{\vec{r}} S_A^{\vec{r}}, \quad (1.5)$$

where $\vec{r}$ labels the superselection sectors set by the electric flux through the boundary links, $p_{\vec{r}}$ is the probability of sector $\vec{r}$, $d_{\vec{r}}$ is the dimension of the associated boundary representation, and $S_A^{\vec{r}}$ is the von Neumann entropy within that sector's block. The first term consists of the Shannon entropy of the sector distribution together with the average $\ln d_{\vec{r}}$ contribution and is therefore purely a gauge contribution. It records which boundary flux configuration is present together with the constraint-locked maximal entanglement of the representation spaces, neither of which is a freely manipulable quantum correlation.

The operational meaning of this split was established by Van Acoleyen et al. and, concurrently, by Soni and Trivedi [45, 46]. Because gauge-invariant operations commute with the sector projectors, no gauge-invariant local protocol can distinguish the pure state from the mixture obtained by depolarizing over the boundary sectors. The entanglement that can be distilled into Bell pairs using gauge-invariant local operations and classical communication is therefore only the second term of Eq. (1.5),

$$E_D(|\psi\rangle) = \sum_{\vec{r}} p_{\vec{r}} S_A^{\vec{r}}, \quad (1.6)$$

the genuinely quantum entanglement residing inside the sectors. The first term is an artifact of embedding the gauge theory into a tensor-product Hilbert space and carries no distillable entanglement.

For Abelian gauge theories at weak coupling, such as the toric code [41], projecting onto a boundary-flux sector leaves a product state across the bipartition. Consequently every $S_A^{\vec{r}}$ vanishes, and the distillable entanglement is exactly zero. The entire entanglement entropy, including its topological correction [47, 48], is therefore contained in the undistillable gauge contribution [45]. Similarly the stabilizer scars of Chapter 3 carry a substantial entanglement entropy yet exactly zero distillable entanglement, and it is the distillable entanglement that exposes their scar structure, where the total entanglement entropy does not.

1.5 Thermalization

As discussed in Sec. 1.1, thermalization in quantum many-body systems is a natural target for quantum simulation. Heavy-ion collisions provide a concrete example: experiments probe the final state of the collision but not the intermediate non-equilibrium dynamics during which the system thermalizes, when quarks and gluons are not directly accessible. Real-time evolution on a quantum simulator could, in principle, access this regime.

Thermalization and irreversibility are also familiar in everyday macroscopic systems: cream stirred into coffee does not spontaneously unmix [49], and a gas released into a vacuum spreads irreversibly throughout its container. Classical statistical mechanics accounts for this behavior through ergodicity, in which chaotic trajectories explore phase space uniformly so that time averages of observables agree with microcanonical averages. At long times, almost all microstates become indistinguishable at the level of macroscopic observables [50].

In contrast, quantum mechanics presents a puzzle. An isolated system prepared in a pure state $|\psi(0)\rangle$ evolves unitarily, $|\psi(t)\rangle = e^{-iHt}|\psi(0)\rangle$, and remains pure for all time. The global von Neumann entropy therefore remains constant, and the dynamics are formally reversible. Nevertheless, many such systems appear to thermalize. This raises the question of how irreversible thermodynamic behavior can emerge from

reversible unitary evolution.

The resolution is that thermalization in a closed quantum system is a statement about *subsystems* and *local observables*, not about the global state. Partition the system into a small region A and its complement $\bar{A}$, and consider the reduced density matrix $\rho_A(t) = \text{Tr}_{\bar{A}} |\psi(t)\rangle\langle\psi(t)|$. The global state stays pure, however the complement acts as an internal reservoir for A , and a system is said to thermalize if $\rho_A(t)$ evolves toward the thermal (Gibbs or microcanonical) reduced state,

$$\lim_{t \rightarrow \infty} \rho_A(t) \approx \rho^{eq} = \frac{1}{Z} e^{-\beta H} \quad (1.7)$$

where β is the inverse temperature for the appropriate energy density. Equivalently, expectation values of local observables, $\langle\psi(t)|\hat{O}|\psi(t)\rangle$, approach their thermal values and exhibit only small fluctuations thereafter. The expression for the thermal value is $\langle\hat{O}\rangle_{thermal} = \frac{1}{\mathcal{N}_{\bar{E}}} \sum_{E_n \in [\bar{E}]} \langle n | \hat{O} | n \rangle$, where $\bar{E}$ is the energy expectation value of the initial state and $\mathcal{N}_{\bar{E}}$ is the number of states in the microcanonical ensemble around $\bar{E}$. The information about the initial state is not destroyed, but it is dispersed into non-local, difficult to measure correlations across the entire system [51].

This leads to the questions of which initial states thermalize and what microscopic mechanism allows unitary evolution to produce this behavior. Expanding the initial state in the energy eigenbasis, $|\psi(0)\rangle = \sum_n c_n |n\rangle$, the expectation value of a local observable as a function of time is written as,

$$\langle\hat{O}(t)\rangle = \sum_n |c_n|^2 O_{nn} + \sum_{n \neq m} c_n^* c_m O_{nm} e^{i(E_n - E_m)t}, \quad (1.8)$$

where $O_{nm} = \langle n | \hat{O} | m \rangle$. If a non-degenerate energy spectrum is assumed, then the off-diagonal terms in Eq.(1.8) cancel for a long-time average, leaving just the diagonal ensemble. For this average to coincide with the microcanonical prediction then the following equation holds,

$$\lim_{t \rightarrow \infty} \frac{1}{T} \int_0^\infty dt \langle\hat{O}(t)\rangle = \sum_n |c_n|^2 O_{nn} = \frac{1}{\mathcal{N}_{\bar{E}}} \sum_{E_n \in [\bar{E}]} \langle n | \hat{O} | n \rangle \quad (1.9)$$

Eq. (1.9) requires that the late-time average expectation value is independent of the fine details of the coefficients c_n , which encode the arbitrary initial condition. Additionally, the individual eigenstate expectation values $\langle n|\hat{O}|n\rangle$ must already vary smoothly with energy and equal the microcanonical value at E_n . These conditions are summarized by the Eigenstate Thermalization Hypothesis (ETH), proposed by Deutsch [52] and Srednicki [53]. This hypothesis postulates that thermalization is encoded in the energy eigenstates themselves, so that a single eigenstate at finite energy density already “looks thermal” to any local probe.

ETH is made precise by expressing the thermalization conditions through an ansatz for the matrix elements of physical observables [51],

$$\langle m|\hat{O}|n\rangle = \bar{O}(E) \delta_{mn} + e^{-S(E)/2} f(E, \omega) R_{mn}, \quad (1.10)$$

where $E = (E_m + E_n)/2$, $\omega = E_m - E_n$, $S(E)$ is the thermodynamic entropy, $\bar{O}(E)$ and $f(E, \omega)$ are smooth functions, and R_{mn} are pseudorandom variables with zero mean and unit variance. The diagonal term reproduces the smooth thermal expectation value. The off-diagonal term is exponentially small in system size and it controls the fluctuations and the relaxation toward equilibrium. Within this framework, an initial state that is a superposition of many *typical* (mid-spectrum) eigenstates dephases, and local observables settle to the value dictated by the smooth function $\bar{O}(\bar{E})$.

ETH is expected to hold for typical states in generic non-integrable, quantum-chaotic Hamiltonians, whose spectral statistics follow the Wigner–Dyson distributions of random matrix theory [51]. Its sharpest structural signature is in entanglement: a typical eigenstate of an ETH-obeying system exhibit *volume-law* entanglement entropy consistent with thermal entropy density, $S_A \propto |A|$, because each subsystem is maximally entangled with the rest as if drawn from a thermal ensemble [50]. ETH does not hold for states in the edges of the spectrum. These states do not have volume-law entanglement due to the density of states being much lower, but instead obey *area-law* entanglement, $S_A \propto |\partial A|$.

Between fully thermal and strongly nonthermal behavior lies the phenomenon of quantum many-body scars (QMBS). A system exhibits QMBS when it is non-integrable and disorder-free, so that the overwhelming majority of its eigenstates obey ETH, yet hosts a small, atypical set of nonthermal eigenstates embedded throughout the many-body spectrum [54, 55, 56, 57, 58, 59]. Because the scarred subspace is vanishingly small relative to the exponentially large Hilbert space, usually growing only polynomially with system size, it does not affect generic thermalization from typical initial states. In an exact scar subspace, the Hamiltonian completely preserves the scar subspace, so it acts as a completely decoupled block of the Hilbert space. Initial states with substantial overlap onto the scar subspace exhibit non-thermal dynamics: rather than exploring the full Hilbert space, its dynamics remain largely confined to the small scar subspace, so its entanglement grows anomalously slowly and its local observables retain memory of the initial condition long after a generic state would have relaxed. Ergodicity is thus broken only *weakly*, on a measure-zero set of states.

The phenomenon was first experimentally observed on a Rydberg-atom quantum simulator [54], with an experiment known as the Rydberg blockade. Rydberg atoms are neutral atoms that have a highly excited state known as a Rydberg state, and they exert Van der Waals forces on each other, making it energetically unfavorable for neighboring atoms to be in the excited state. When the system is prepared in the Néel state, $|\mathbb{Z}_2\rangle = |grgr\dots\rangle$ ($|g\rangle$ is the ground state and $|r\rangle$ is the Rydberg state), and then evolved, the only energetically favorable state to evolve into is $|\mathbb{Z}'_2\rangle = |rgrg\dots\rangle$. The evolution between these two states displayed persistent, slowly decaying oscillations of local observables rather than the rapid relaxation expected by ETH of a non-integrable model, as depicted in Fig. 1.1(c). Turner *et al.* [60] traced this anomaly to a band of atypical, low-entanglement, equally spaced, energy eigenstates threading the otherwise chaotic spectrum of the PXP model, which are depicted in Fig. 1.1(a) and (b). The PXP model encodes the same blockade constraint as the Rydberg exper-




Figure 1.1: *QMBS in the Rydberg Blockade and PXP Model*. (a) depicts the energy spectrum of an example non-integrable system. The dense region of the spectrum is comprised of the typical states, and the top and bottom sparse regions are the groundstate-like states. (b) is a plot from [60] and depicts the von Neumann entropy of the all the eigenstate of the PXP model. The x-axis is the energy of the eigenstate and the y-axis in the von Neumann entropy. The numbered states are the sub volume-law entanglement QMBS states and they correspond to the red lines in the energy spectrum in (a). The color corresponds to the density of states. (c) is a plot from [54] that depicts the oscillation of of a local observable as a function of time for the Néel state in the Rydberg blockade experiment. The y-axis is the domain-wall density, which corresponds to the density of number of neighboring atoms that are in the same state. This oscillation is also known as a revival. A non-scar would relax in about one coupling time (one Rabi or oscillation period), whereas the scar dynamics stay coherent for many such periods. This plot includes data from a 9 atom chain, a 51 atom chain, and a matrix product state (MPS) simulation.

iment and captures the constrained structure of the accessible Hilbert space in which scarring arises [61, 62, 63, 64, 65]. The equally spaced energy eigenstates produce long-lived period revivals that match the oscillations observed in the Rydberg-atom experiment. The QMBS terminology borrows from single-particle quantum scars [66], where individual eigenfunctions of a classically chaotic billiard exhibit anomalously enhanced amplitude along unstable periodic orbits. The quantum many-body scar analogues are thin, nonthermal filaments running through an otherwise ergodic spectrum. Their defining hallmarks are sub-thermal, typically logarithmic or area-law, entanglement entropy at energies where ETH would mandate a volume law, and persistently nonthermal dynamics from special initial states that fail to relax. Many constructions of QMBS have equally spaced energy eigenstates that produce period revivals for observable expectation values.

Many scar families admit exact analytic construction and can be viewed as integrable structures within non-integrable systems. Towers of equally spaced scar states can be generated by spectrum-generating algebras [67, 68, 69, 70, 71, 72, 73, 74, 75], as in the exact excited states of the AKLT chain [76, 77, 74]; they can be engineered by embedding a nonthermal subspace into a thermal Hamiltonian [78, 79, 80, 81] or expressed through matrix-product-state constructions. Quantum many-body scars are not confined to engineered spin chains: they have been found in lattice gauge theories, including the Schwinger model, $U(1)$ quantum link models, and non-Abelian formulations [82, 83, 84, 85, 84, 86, 82]. The Rydberg blockade and the Gauss law of a gauge theory are the same kind of object; a local constraint restricting the dynamics to a sub-sector of the Hilbert space, and the PXP model is itself dual to a $U(1)$ lattice gauge theory [87], so it is natural that scarring recurs in the gauge theory setting. This places weak ergodicity breaking squarely within the gauge theory framework that motivates much of quantum simulation for high-energy and nuclear physics, and makes the thermalization properties of gauge theories a question of direct experimental relevance. The mechanisms of all families of scars are unified

by the language of commutant algebras [88, 89, 90, 91, 92], in which the scarred, or more generally fragmented, subspace is identified with a nontrivial commutant of the algebra generated by the Hamiltonian’s local terms. The existence of conserved structure beyond the model’s manifest symmetries signals a protected, dynamically decoupled subspace. This commutant framework will be used in Chapter 3 to derive scar solutions analytically.

1.6 Computational Complexity and Nonstabilizerness

The previous sections classified eigenstates primarily by their entanglement structure, distinguishing volume-law behavior in the thermal bulk from area-law states near spectral edges. However, entanglement alone is not a reliable indicator of classical computational complexity. A complementary notion is *nonstabilizerness*, or “magic”, which quantifies the extent to which a quantum state lies outside the class of efficiently classically simulable stabilizer states. Importantly, highly entangled states can still be efficiently representable on a classical computer, making entanglement an incomplete proxy for computational hardness [93].

This distinction is formalized by the stabilizer framework. Let $\mathcal{P}_n$ denote the n -qubit Pauli group, generated by tensor products of the single-qubit Pauli operators $\{I, X, Y, Z\}$ together with the phases $\{\pm 1, \pm i\}$. A *stabilizer state* $|\psi\rangle$ is the unique simultaneous $+1$ eigenstate of a stabilizer group $\mathcal{S} \subset \mathcal{P}_n$: an abelian subgroup generated by n independent commuting Pauli operators and not containing $-I$, so that $\mathcal{S}$ has 2^n elements and fixes $|\psi\rangle$ completely. Stabilizer states are precisely those obtainable from $|0\rangle^{\otimes n}$ by Clifford unitaries, which map Pauli operators to Pauli operators under conjugation and are generated by the Hadamard, phase, and controlled-NOT gates $\{H, S, \text{CNOT}\}$ [9]. Crucially, many highly entangled states are stabilizer states: GHZ and cat states, graph and cluster states, and the toric-code ground state, which is long-range (topologically) entangled [94, 41].

The distinction between entanglement and classical hardness is made precise by

the Gottesman–Knill theorem [93, 95]: any quantum circuit composed of stabilizer-state inputs, Clifford gates, and measurements of Pauli observables can be simulated efficiently on a classical computer. The efficiency arises from the compact representation of stabilizer states. Instead of storing 2^n amplitudes, one tracks the n stabilizer generators, each an n -qubit Pauli string, which can be updated under Clifford gates in polynomial time. Thus, entanglement alone is insufficient for quantum computational advantage: even states with long-range entanglement, such as the toric code ground state, remain classically simulable within this framework.

To achieve universal quantum computation, one must go beyond the Clifford-stabilizer framework. Supplementing Clifford circuits with even a single non-Clifford gate, such as the $T = \text{diag}(1, e^{i\pi/4})$ gate, or, equivalently, with a supply of non-stabilizer “magic states” to be consumed by gate teleportation, yields universality [96, 97]. The consumption of such resources defines a quantitative notion of *magic*, or nonstabilizerness. Magic is a kind of resource whose free states are the stabilizer states and whose free operations are the Clifford operations. Together with entanglement, it gauges a state’s distance from classical tractability. As noted, a state which is both modestly entangled and a stabilizer state is classically simulable, and universal, classically hard computation demands both resources in extensive supply.

Several measures of magic exist, but the stabilizer Rényi entropy (SRE) [98] has become the standard many-body diagnostic because it is computable directly from Pauli expectation values. For a pure state $|\psi\rangle$ of n qubits with Hilbert-space dimension $d = 2^n$, the squared Pauli expectation values define, after normalization, a probability distribution over the Pauli group,

$$\xi_{\hat{P}} = \frac{1}{d} \langle \psi | \hat{P} | \psi \rangle^2, \quad \hat{P} \in \mathcal{P}_n, \quad \sum_{\hat{P} \in \mathcal{P}_n} \xi_{\hat{P}} = 1, \quad (1.11)$$

where the normalization uses $\sum_{\hat{P}} \langle \psi | \hat{P} | \psi \rangle^2 = d$ for any pure state. The α -SRE is

essentially the Rényi- α entropy of this distribution,

$$M_\alpha(|\psi\rangle) = \frac{1}{1-\alpha} \log \left(\frac{1}{d} \sum_{\hat{P} \in \mathcal{P}_n} \langle \psi | \hat{P} | \psi \rangle^{2\alpha} \right). \quad (1.12)$$

A stabilizer state has $\langle \psi | \hat{P} | \psi \rangle = \pm 1$ on the d elements of its stabilizer group and 0 on every other Pauli operators, so $\xi_{\hat{P}}$ is uniform over d outcomes and $M_\alpha = 0$. By definition stabilizer states carry exactly zero magic, and the SRE agrees. Generic many-body states exhibit extensive magic, with SRE growing proportionally to system size in chaotic regimes [99, 100], placing them firmly in the classically intractable regime. Nonstabilizerness has likewise been studied in kinetically constrained, scarred systems. In the PXP model the Rydberg blockade generates many-body magic and the scarred eigenstates are themselves strongly non-stabilizer [101].

Viewing ergodicity-breaking states through the complementary lens of computational resources frames the central contribution of Chapter 3. There, we identify quantum many-body scars in a $\mathbb{Z}_2$ lattice gauge theory whose scar subspace is spanned entirely by *stabilizer* states, states of exactly zero magic in the sense of Eq. (1.12). What distinguishes these scars is not just anomalously low entanglement entropy, but both their distillable entanglement and their magic are exactly zero, which exposes a direct link between quantum ergodicity and computational complexity. Additionally, the scars form an exact, dynamically decoupled subspace whose dimension grows only linearly with the system size, so that a state initialized within it remains there and evolves under a small effective Hamiltonian. As a result, real-time evolution within this subspace is efficiently classically simulable. Since this work has been published, more families of stabilizer scars have been identified [102, 103]. In Chapter 4, the stabilizer character of this subspace is used to bound the cost of fidelity estimation, and allows the scars to be repurposed from objects of study into a tool for benchmarking quantum simulations at scale.

1.7 Quantum Verification

Reducing the cost of a simulation is only part of the challenge, since a computation is useful only if its output can be trusted. This is particularly difficult in precisely the regime where quantum simulation is most valuable: when classical simulation is impossible. In this setting there is no classical reference for comparison, and near-term devices further introduce noise and control errors. Certifying that a quantum processor has implemented the intended state or dynamics is therefore as fundamental as reducing computational cost [104, 31].

A range of certification tools exists, each with inherent limitations. Full quantum state tomography [105, 106] is informationally complete but requires resources that scale exponentially with system size. Randomized benchmarking [107, 108, 109] characterizes average gate performance rather than the fidelity of a specific state or circuit output. Cross-entropy benchmarking, used in demonstrations of quantum advantage in random circuit sampling [30], certifies sampling distributions rather than physical observables of a target state.

Direct fidelity estimation [110, 111] provides a more relevant primitive, estimating the overlap between a prepared state and a target state by sampling Pauli observables. Its efficiency depends on the distribution of the target state’s weight in the Pauli basis; states whose Pauli support is concentrated require only a modest number of measurements, while broadly distributed states require exponentially many samples. However, this approach implicitly assumes that the target state is classically tractable, since its Pauli expectation values must be efficiently computable.

Both requirements, a classically computable target and a manageable measurement cost, are governed by the notion of nonstabilizerness introduced in Sec. 1.6. States with low magic are efficiently describable classically and have concentrated weight in the Pauli basis, making them amenable to efficient certification. By contrast, highly magic states are neither efficiently representable nor efficiently certifiable

in this framework [112]. Thus, low-nonstabilizerness states occupy a regime that is simultaneously classically accessible and experimentally verifiable.

The remaining challenge is to benchmark quantum processors precisely in the regime where classical prediction is impossible. Chapter 4 resolves this by embedding a classically solvable reference inside the intractable problem. The stabilizer scars of Chapter 3 evolve on the same device and under the same Hamiltonian as the thermalizing bulk, yet their dynamics can be computed exactly and certified cheaply, so that the fidelity measured on the scar states bounds the fidelity of the genuinely intractable, non-scar dynamics running alongside them. Structure identified in one part of the system thus becomes a tool for certifying dynamics that have no classical reference.

Chapter 2

FASTER SPECTRAL DENSITY CALCULATION USING ENERGY MOMENTS

This work has been published as “Faster spectral density calculation using energy moments” by Jeremy Hartse and Alessandro Roggero, in The European Physical Journal A [113]

2.1 Introduction

As discussed in Sec. 1.3, response functions describe the linear response of a many-body system after an excitation and contain the same information as an inclusive reaction cross section. They can be expressed in terms of the spectral density operator $\delta(\hat{H} - \omega)$ of the Hamiltonian $\hat{H}$ describing the many-body system. First principle calculations of response function is in general extremely challenging for strongly correlated systems. A very powerful approach to study dynamical properties of many-body systems is to employ integral transform techniques which map the local spectral density into more manageable ground state expectation value which can be used to infer properties of the response function. This is the approach used in the Lorentz Integral Transform (LIT) method [37, 114] and the more recent Gaussian Integral Transform (GIT) method [115, 116]. Thanks to the ability to efficiently simulate the real time dynamics of many-body systems, simulations employing quantum computers offer the possibility to tackle the calculation of scattering cross sections from first principles (see e.g. [11] for a recent review). Interestingly, in order to describe inclusive scattering in the linear response regime using the response function, integral transform techniques are also useful in the design of efficient quantum algorithms [15, 115]. It is

therefore important and timely to extend the available techniques in order to reduce as much as possible the computational resources required to calculate the nuclear response function on a quantum device.

Given an Hamiltonian $\hat{H}$, an initial state $|\Psi_0\rangle$ and a Hermitian excitation operator $\hat{O}$, our goal is to evaluate the frequency dependent response function defined as

$$S(\omega) = \langle \Psi_0 | \hat{O}^\dagger \delta(\omega - \hat{H}) \hat{O} | \Psi_0 \rangle = \sum_m \left| \langle \Psi_0 | \hat{O} | m \rangle \right|^2 \delta(\omega - E_n) \quad (2.1)$$

where in the last line we used the expansion on the eigenbasis $\{|m\rangle\}$ of $\hat{H}$. In general, it is difficult to directly use this definition as it requires a complete knowledge of the full energy eigenspectrum. The main idea is, similarly to the LIT and the GIT methods, to consider instead an integral transform with kernel K given by

$$\begin{aligned} \Phi(\nu) &= \int d\omega K(\nu, \omega) S(\omega) \\ &= \langle \Psi_0 | \hat{O}^\dagger K(\nu, \hat{H}) \hat{O} | \Psi_0 \rangle \\ &= \sum_n \left| \langle \Psi_0 | \hat{O} | n \rangle \right|^2 K(\nu, E_n) . \end{aligned} \quad (2.2)$$

Note that both the original transform and its integral transform have units on inverse energy. Furthermore, in this work we will focus on translationally invariant kernel functions for which $K(\nu, \omega) = K(\nu - \omega)$, but extensions to the more general case are straightforward.

Once $\Phi(\nu)$ has been obtained, one usually attempts an inversion of the integral transform in order to obtain $S(\omega)$ back, however this procedure can introduce uncontrollable errors whenever the kernel function has compact support [117, 118]. A variety of approximate inversion techniques that introduce, more or less explicitly, additional smoothing to reduce these artifacts have been proposed in the past [119, 120, 121, 122, 123]. However, if the kernel function is chosen appropriately this last step might not be necessary [115, 116]. In particular, it is convenient to consider kernels such that

$$\int_{-\Delta}^{\Delta} d\nu K(\nu) \geq 1 - \Sigma , \quad (2.3)$$

A kernel satisfying Eq. (2.3) is called Σ -*approximate* with *resolution* Δ [115]. The reason for this definition is that, in the commonly encountered situation where one is interested in observables of the form

$$Q(S, f) = \int d\omega S(\omega) f(\omega) , \quad (2.4)$$

for some bounded function f , one can use directly the integral transform Φ to approximate $Q(S, f)$ with $Q(\Phi, f)$ with controllable error [115]. More intuitively, we can think of Φ as a finite width approximation of the original response function with an energy resolution given by Δ and additional tails controlled by the parameter Σ . The original response is then recovered in the limit $\Delta, \Sigma \rightarrow 0$.

This approach was recently used in Ref. [116] to provide a controllable approximation of the spectral density using histograms derived from the integral response (see also [124] for a recent application for computing spectral functions in light nuclei using the Coupled Cluster method).

A direct procedure to construct integral kernels satisfying Eq. (2.3) is to use a polynomial expansion such as

$$K(\nu - \omega) = \sum_n^{\infty} c_n(\nu) \phi_n(\omega) , \quad (2.5)$$

with $\{\phi_n\}$ an orthonormal basis of polynomials. This leads directly to an alternative expression for the integral transform in this basis

$$\Phi(\nu) = \sum_n^{\infty} c_n(\nu) \langle \Psi_0 | \hat{O}^\dagger \phi_n(\hat{H}) \hat{O} | \Psi_0 \rangle = \sum_n^{\infty} c_n(\nu) m_n , \quad (2.6)$$

where m_n denotes the frequency moments of the response function over these polynomials

$$m_n = \int d\omega \phi_n(\omega) S(\omega) . \quad (2.7)$$

If the series in Eq. (2.6) converges rapidly, one can then estimate the integral transform $\Phi(\nu)$ with a small error by keeping only the first N terms in the expansion. Since

the response function, and thus its integral transform, have units of inverse energy we will quantify the error in its approximation as ε/Ω with $\varepsilon > 0$ and Ω a suitable energy constant. As shown in [115], we can approximate a response function with a Σ -accurate kernel with resolution Δ using a basis of Chebyshev polynomials and a number of terms given by

$$N \approx \tilde{\mathcal{O}} \left(\frac{\|\hat{H}\|}{\Delta} \sqrt{\log \left(\frac{1}{\Sigma} \right) \log \left(\frac{\Omega}{\Delta \varepsilon} \right)} \right) , \quad (2.8)$$

with $\|\hat{H}\|$ as the spectral norm of the Hamiltonian. The notation used here neglects subleading logarithmic factors, in particular we define $\tilde{\mathcal{O}}(f(x)) = \mathcal{O}(f(x) \log(f(x)))$ in line with previously reported estimates [115].

This results applies to any possible response function and is helpful whenever the Chebyshev moments of the Hamiltonian can be computed efficiently. This can be done in practice with quantum computers by means of a suitable quantum walk [125, 34, 126, 127, 128] or in approximate way with classical methods, such as the Coupled Cluster approach [116, 124]. In many applications the spectral function has a number of distinctive features like being dominated by low energy contributions or displaying a distinct peak structure like e.g. in the quasi-elastic regime. Many of these features are captured by energy moments of the form [129]

$$\mu_n = \int d\omega \omega^n S(\omega) . \quad (2.9)$$

For instance the qualitative shape of a quasi-elastic peak can be characterized with a good accuracy from the first few moments alone [130]. The advantage of using moments μ_n to characterize the spectral density is that in many situations they can be calculated explicitly using ground-state many-body methods (e.g. with Monte Carlo [131, 132, 133]).

In this work we employ a Fourier basis for the polynomial expansion of the integral transform Eq. (2.6) in order to incorporate this information into the calculation of the full spectral density and reduce the overall computational cost of the simulation.

Thanks to their ability to efficiently simulate the real-time evolution of many-body systems, quantum computers are expected to be able to give us access to expectation values of the form

$$g_\psi(t) = \langle \psi | e^{it\hat{H}} | \psi \rangle, \quad (2.10)$$

for states ψ which can be easily prepared. The function $g_\psi(t)$ contains important information about the many body system, and several authors have proposed techniques to use this to get access to a variety of observables like excitation energies, the local spectral density and even thermodynamic expectation values [134, 135, 136, 137, 138].

Under general consideration we can expect two main time scales to play a role: first, in order to obtain a frequency resolution of order Δ , the maximum time required will need to scale as $T = \mathcal{O}(1/\Delta)$; second, since the full energy spectrum is contained in a frequency interval of size $\|\hat{H}\|$, a time step $\delta t = \mathcal{O}(1/\|\hat{H}\|)$ will guarantee a perfect reconstruction without aliasing thanks to the Shannon-Nyquist theorem. The combination of these two time scales predicts a scaling of the number of time-steps required as $T/\delta t = \mathcal{O}(\|\hat{H}\|/\Delta)$. However, as shown recently in Ref. [136], in situations where the spectral function has an energy variance $\sigma^2 \ll \|\hat{H}\|^2$, a good Fourier approximation can be obtained with a scaling $T/\delta t = \mathcal{O}(\sigma/\Delta)$ instead. The goal of this work is to formalize this intuition and provide rigorous error bounds allowing use of prior information about the energy moments of the spectral function in order to reduce the number of expectation values required.

In Section 2.2 we present the general framework employing Fourier moments to approximate the response function and in specialize the treatment to the Gaussian Integral Transform in Section 2.3. Section 2.4 details two example response functions that are reconstructed with this method and the improved efficiency afforded by employing energy moments. In Section 2.5 we conclude by discussing applications of this method to typical scattering properties in nuclear physics, and the estimated improvement in number of Fourier moments required.

2.2 Fourier Based Reconstruction

In order to use directly a discrete Fourier transform to perform the expansion of the integral kernel in Eq. (2.5) it's convenient to define a periodic extension of the kernel function as follows [135]

$$K^\chi(\nu, \omega) = \sum_{n=-\infty}^{\infty} K\left(\nu, \omega + n\chi\|\hat{H}\|\right), \quad (2.11)$$

where the period $P = \chi\|\hat{H}\|$ is expressed here as a function of the spectral norm. This convention allows us to truncate the sum at the first term if we take $\chi \gg 1$: since the spectrum is bounded only the term with $n = 0$ has a finite contribution. As χ is reduced more terms will start to contribute to the sum and the extended integral kernel K^χ will start to deviate from the original. The main advantage of this construction is that we can express directly the periodically-extended kernel as a Fourier series

$$K^\chi(\nu, \omega) = \frac{1}{\chi\|\hat{H}\|} \sum_{n=-\infty}^{\infty} k_n^\chi(\nu) \exp\left(-i \frac{2\pi n}{\chi\|\hat{H}\|} \omega\right), \quad (2.12)$$

where the Fourier transformed coefficients are given by

$$k_n^\chi(\nu) = \int_{-\infty}^{\infty} d\omega \exp\left(i \frac{2\pi n}{\chi\|\hat{H}\|} \omega\right) K(\nu, \omega). \quad (2.13)$$

The main advantage of working with K^χ and this Fourier series representation is that we can now express a general integral transform of the response function as a linear combination of expectation values of the time-evolution operator. From the definition in Eq. (2.2) we can in fact express the new integral transform as follows

$$\begin{aligned} \Phi^\chi(\nu) &= \langle \Psi_0 | \hat{O}^\dagger K^\chi(\nu, \hat{H}) \hat{O} | \Psi_0 \rangle \\ &= \frac{1}{\chi\|\hat{H}\|} \sum_{n=-\infty}^{\infty} k_n^\chi(\nu) \langle \Psi_0 | \hat{O}^\dagger e^{-in\delta t \hat{H}} \hat{O} | \Psi_0 \rangle \end{aligned} \quad (2.14)$$

with finite time-steps of size $\delta t = 2\pi/(\chi\|\hat{H}\|)$. Provided we choose a smooth integral kernel in the ω variable, the series expansion converges quickly and we can therefore

obtain an accurate approximation of the integral transform by taking a truncation to some finite order

$$\Phi_N^\chi(\nu) = \frac{1}{\chi \|\hat{H}\|} \sum_{n=-N}^N k_n^\chi(\nu) \langle \Psi_0 | \hat{O}^\dagger e^{-in\delta t \hat{H}} \hat{O} | \Psi_0 \rangle . \quad (2.15)$$

For example, if we consider integral kernels in C^∞ , the convergence will be in general super-polynomial in N and, importantly, rigorous bounds on the truncation error can be found with relatively straightforward calculations. It is important to note at this point that the use of a smoothing kernel with a finite energy resolution Δ is critical to ensure that the expansion is well-behaved: indeed a direct approximation of the response function $S(\omega)$ using a finite Fourier sum would have encountered difficulties due to the discreteness of the energy spectrum for a finite system, which creates sharp peaks. The difference between the present approach and related ones (from Refs. [115, 116]) with more heuristic smoothing procedures like the Maximum Entropy Method [119, 121] or alternative polynomial expansion methods like the Kernel Polynomial Method [139] is that the energy window over which the smoothing is applied is fully under control and all errors can be accounted for. The price to pay to be able to use higher energy resolution is a corresponding increase in the number of terms in the expansion Eq. (2.15) which corresponds to a comparable increase in the maximum time $T = N\delta t$ over which one needs to be able to simulate the many-body dynamics. We will discuss this in more detail for a specific choice of integral kernel in the next section.

2.3 *Gaussian Integral Transform*

As shown already in Ref. [115], the Gaussian Integral Transform (GIT) is particularly useful for the purpose of obtaining a Σ -accurate integral transform with resolution Δ with a fast converging polynomial expansion like Eq. (2.5). This result is rather intuitive since a Gaussian envelope is a perfect trade-off between good resolution in frequency and small widths in the time domain. The original construction from

Ref. [115] used a Chebyshev expansion for the polynomial basis, however in this work we instead explore an expansion of the Gaussian kernel into Fourier modes.

The first step is to determine an appropriate value for the width Λ of the Gaussian kernel in order to satisfy the condition in Eq. (2.3). A direct calculation gives

$$\frac{1}{\sqrt{2\pi}\Lambda} \int_{-\Delta}^{\Delta} d\nu \exp\left(-\frac{\nu^2}{2\Lambda^2}\right) = \operatorname{erf}\left(\frac{\Delta}{\sqrt{2}\Lambda}\right) \geq 1 - \Sigma. \quad (2.16)$$

Using the upper-bound on the complementary error function $\operatorname{erfc}(x) = 1 - \operatorname{erf}(x)$ by a Gaussian

$$\operatorname{erfc}(x) \leq \exp(-x^2), \quad (2.17)$$

we can find the following sufficient condition [115]

$$\Sigma \geq \exp\left(-\frac{\Delta^2}{2\Lambda^2}\right) \Rightarrow \Lambda \leq \frac{\Delta}{\sqrt{2\log(1/\Sigma)}}. \quad (2.18)$$

Following the notation from Eq. (2.11) and Eq. (2.12), we can then express the periodically extended Gaussian kernel as

$$G^x(\nu, \omega) = \frac{1}{\chi\|\hat{H}\|} \sum_{n=-\infty}^{\infty} g_n^x(\nu) \exp(-i\delta t n \omega) \quad (2.19)$$

with the corresponding Fourier coefficients

$$g_n^x(\nu) = \exp(-i\delta t n \nu) \exp\left(-\frac{\delta t^2 \Lambda^2}{2} n^2\right). \quad (2.20)$$

In the expressions above we used the definition $\delta t = 2\pi/(\chi\|\hat{H}\|)$ for the time-step. We can now write the approximate integral transform obtained by truncating the series as

$$\Phi_N^x(\nu) = \frac{1}{\chi\|\hat{H}\|} \sum_{n=-N}^N g_n^x(\nu) \langle \Psi_0 | \hat{O}^\dagger e^{-in\delta t \hat{H}} \hat{O} | \Psi_0 \rangle. \quad (2.21)$$

The ideal integral transform $\Phi(\nu)$ would be obtained in principle by choosing Λ in order to satisfy Eq. (2.18), taking $\chi = 1$ and letting $N \rightarrow \infty$. The error in the approximation Φ_N^x is caused only by taking a finite number of terms N and (possibly)

reducing the value of χ which parametrizes the frequency interval used in the periodic extension. For a fixed value of the frequency ν we can then bound the error with

$$|\Phi_N^\chi(\nu) - \Phi(\nu)| \leq |\Phi^\chi(\nu) - \Phi(\nu)| + |\Phi_N^\chi(\nu) - \Phi^\chi(\nu)| , \quad (2.22)$$

where we simply used the triangle inequality. The advantage is that the two error contributions on the right hand side can be bounded individually in a simpler way. We will denote the first error on the right hand side as ϵ_P and the second as ϵ_N . In addition to these two sources of error, we need also to account for the fact we will estimate the Fourier moment

$$m_n^\chi = \langle \Psi_0 | \hat{O}^\dagger e^{-in\delta t \hat{H}} \hat{O} | \Psi_0 \rangle \quad (2.23)$$

by computing the expectation value with a finite statistical sample leading to an additional error

$$\epsilon_S = |\Phi_N^\chi(\nu) - \bar{\Phi}_N^\chi(\nu)| , \quad (2.24)$$

where we have denoted by $\bar{\Phi}_N^\chi(\nu)$ the finite sample estimate of $\Phi_N^\chi(\nu)$. The total error will be given then by the sum of all contributions $\epsilon = \epsilon_P + \epsilon_N + \epsilon_S$.

Before presenting our results for the complexity of estimating an accurate approximation of the integral transform $\Phi(\nu)$ it is convenient to specify several of the conventions we use, note however that the results provided in App. 2.A are completely general and do not depend on these conventions. First of all, since the response function (and thus it's integral transform) have dimensions of inverse energy, we will consider a dimensionless error parameter $\varepsilon > 0$ obtained using a suitable energy scale Ω . That is, we want to find values χ and N for which

$$|\bar{\Phi}_N^\chi(\nu) - \Phi(\nu)| \leq \frac{\varepsilon}{\Omega} . \quad (2.25)$$

Second, we consider the situation where we are interested in approximating the response function on a finite energy window $\omega \in [\omega_{\min}, \omega_{\max}]$. For many applications in nuclear physics is also customary to shift the Hamiltonian by the ground

state energy so that all frequencies become positive. Since both of these considerations apply directly to the integral transform, we will consider the situation where $\hat{H}$ has been shifted so that the ground state is at $\omega_{\min} = -\|\hat{H}\|$ consider a range $[-\|\hat{H}\|, -\|\hat{H}\| + \delta\nu]$ for the energies we want the value of $\Phi(\nu)$. Finally, since we are interested in computing the Fourier moments m_n^χ on a quantum computer, we will consider the situation where the excitation operator has been appropriately rescaled so that the state $\hat{O}|\Psi_0\rangle$ is normalized to one. This directly implies that the zeroth moment μ_0 will also be equal to one. A rescaling of this form has been employed already in past works on quantum algorithms for the response function [15, 115, 38, 116, 140] and is natural when using efficient methods for preparing the excited state (see [16]).

The only error term that depends directly on the energy variance $\sigma^2 = \mu_2 - \mu_1^2$ is the first one while the other one will only depend parametrically on the chosen period $P = \chi\|\hat{H}\|$. We start with the truncation error which can be bounded easily using standard techniques resulting in the following requirement for the number of terms

$$N \geq \frac{\chi\|\hat{H}\|}{\sqrt{2\pi}\Lambda} \sqrt{\log\left(0.4\frac{\Omega}{\varepsilon_N\Lambda}\right)}. \quad (2.26)$$

A full derivation of this result can be found in App. 2.A.3. The dependence on the error is typical for Gaussian kernels (c.f. [115, 116]). The statistical contribution of the error can be controlled, with confidence level $1 - \delta$, by requiring in the worst case a number of experiments given by

$$S = N \frac{\Omega^2}{\varepsilon_S^2 \Lambda^2} \log\left(\frac{2}{\delta}\right) = \frac{\chi\|\hat{H}\|\Omega^2}{\sqrt{2\pi}\varepsilon_S^2 \Lambda^3} \sqrt{\log\left(0.4\frac{\Omega}{\varepsilon_N\Lambda}\right)} \log\left(\frac{2}{\delta}\right), \quad (2.27)$$

where in the second line we used the estimate from Eq. (2.26). We present a full derivation of this result in App. 2.A.4. As we can see, the sample complexity of the scheme is directly proportional to the factor χ controlling the period.

The remaining error ϵ_P , controlled by the choice of χ , can be found under to separate situations: the general case where we do not use information from known

energy moments μ_n and the typical case where we have information at least on the energy variance. As mentioned in the text following Eq. (2.11), in the first situation we want to take $\chi \gg 1$ in order to minimize the error. In particular, as shown explicitly in App. 2.A.1, we find that for $\epsilon_P \leq \epsilon_P/\Omega$ the following choice would be sufficient

$$\chi = 2 + \frac{\sqrt{2}\Lambda}{\|\hat{H}\|} \sqrt{\log \left(\frac{2\Omega}{\epsilon_P \Lambda} \right)}. \quad (2.28)$$

In the limit $\Lambda \rightarrow 0$ we recover the Shannon-Nyquist theorem which gives $\chi = 2$ for a perfect reconstruction. At this point we can look for the asymptotic scaling of both N and S while guaranteeing a total error $\epsilon_P + \epsilon_N + \epsilon_S \leq \epsilon$. If take $\epsilon_P = \epsilon_N = \epsilon_S = \epsilon/3$ we find immediately

$$N = \tilde{\mathcal{O}} \left(\frac{\|\hat{H}\|}{\Delta} \sqrt{\log \left(\frac{1}{\Sigma} \right) \log \left(\frac{\Omega}{\epsilon \Delta} \right)} + \log \left(\frac{\Omega}{\epsilon \Delta} \sqrt{\log \left(\frac{1}{\Sigma} \right)} \right) \right), \quad (2.29)$$

which, apart from a logarithmic correction, is the same result found for the Chebyshev version of the GIT protocol. In turn we find that the required number of samples will scale in general as

$$S = \tilde{\mathcal{O}} \left(\frac{\|\hat{H}\| \Omega^2}{\epsilon^2 \Delta^3} \log^{3/2} \left(\frac{1}{\Sigma} \right) \log \left(\frac{1}{\delta} \right) \right) \quad (2.30)$$

again similar to the Chebyshev version (c.f. App. 2.A.4 and the original work Ref. [115]). The main difference is that the Chebyshev implementation avoids the (necessary) approximation of the time-evolution operator. If simulations schemes with optimal asymptotic scaling with the error are used, like the one based on qubitization [34, 126], the Chebyshev GIT will have an advantage in terms of gate counts over its Fourier version described here.

We now turn to the main result of this work where instead we use information about first two energy moments to design a scheme with better scaling. The main result we use is the Chebyshev inequality which, in terms of response functions, states that (assuming $\mu_0 = 1$)

$$1 - \int_{\mu_1 - \Gamma}^{\mu_1 + \Gamma} d\omega S(\omega) = \text{Prob} [|\omega - \mu_1| \geq \Gamma] \leq \frac{\sigma^2}{\Gamma^2}, \quad (2.31)$$

for any positive constant energy Γ . The idea is to use Eq. (2.31) to constrain the integrated strength of the spectrum away from the mean. The interested reader can find the full derivation in App. 2.A.2, the final result is

$$\chi = \frac{2.7}{\varepsilon_P^{1/3}} \frac{\Omega^{1/3} \sigma^{2/3}}{\|\hat{H}\|} + \frac{\delta\nu}{\|\hat{H}\|}, \quad (2.32)$$

which is valid for $\Lambda \leq 2\sigma$ (see App. 2.A.2 for an estimate valid also in a lower resolution regime). This immediately gives an estimate for the required number of moments

$$N = \mathcal{O} \left(\frac{1}{\Delta} \left(\frac{\Omega^{1/3} \sigma^{2/3}}{\varepsilon^{1/3}} + \delta\nu \right) \sqrt{\log \left(\frac{1}{\Sigma} \right) \log \left(\frac{\Omega}{\varepsilon \Delta} \right)} \right) \quad (2.33)$$

Several comments are in order at this point. First of all, the scaling of N with the target error ε is exponentially worse than the original result Eq. (2.29) which didn't use energy moments. For situations that require very small errors $\varepsilon \lesssim \Omega \sigma^2 / \|\hat{H}\|^3$ then the general scheme with χ from Eq. (2.28) should be used instead. This drawback can be mitigated if more information is available. For instance if we know the value of the n central moment $\widetilde{\mu}_n$ then we can bring the cost down to

$$N = \mathcal{O} \left(\frac{\Omega^{1/(n+1)} \widetilde{\mu}_n^{1/(n+1)}}{\Delta \varepsilon^{1/(n+1)}} \sqrt{\log \left(\frac{1}{\Sigma} \right) \log \left(\frac{\Omega}{\varepsilon \Delta} \right)} + \frac{\delta\nu}{\Delta} \sqrt{\log \left(\frac{1}{\Sigma} \right) \log \left(\frac{\Omega}{\varepsilon \Delta} \right)} \right). \quad (2.34)$$

See App. 2.A.2 for additional details and a proof.

The second important comment we can make about the result Eq. (2.33), which also applies to its generalization Eq. (2.34), is that if we want to reconstruct the energy spectrum over a large frequency range scaling as $\mathcal{O}(\|\hat{H}\|)$ then the second contribution proportional to $\delta\nu$ will dominate and we recover again the general result obtained without moments. This somewhat counter-intuitive behavior can be understood by noticing that our error metric in Eq. (2.25) measures the error in a pointwise fashion: a very narrow but tall peak in the energy spectrum will contribute to the error even

if its spectral weight is small. For specific applications, like the calculation of energy histograms presented in Ref. [116], better bounds could be obtained. We leave this development for future work.

Finally we note that all the results presented in this work depend on an arbitrary energy scale Ω which controls the energy dimension of the final error. An appropriate choice for this parameter depends on the specific physical application desired. For example, in applications where one is interested in obtaining histograms of the spectrum like in Ref. [116], the choice $\Omega = \Delta$ seems appropriate.

2.4 Numerical Examples

In this section we present some numerical experiments that show in practice the savings afforded by the method present here. In order to show the general trend, we will employ two simple model response functions $S_A(\omega)$ and $S_B(\omega)$. In both cases we produce a peak at low frequencies according to a skewed Gaussian distribution

$$S_{peak}(\omega) = \frac{1}{\beta\sqrt{2\pi}} \exp\left\{-\frac{(\omega - \xi)^2}{2\beta^2}\right\} \left(1 + \operatorname{erf}\left(\frac{\alpha(\omega - \xi)}{\sqrt{2}\beta}\right)\right), \quad (2.35)$$

where ξ is the location of the peak, β is the scale of the distribution and α the skewedness. In order to explore the impact of a tail at high energies we also use

$$S_{tail}(\omega) = \begin{cases} 0 & \text{for } \omega < \omega_{thr} \\ \lambda\rho(|\omega - \omega_{thr}|^\gamma + \rho)^{-1} & \text{for } \omega \geq \omega_{thr} \end{cases} \quad (2.36)$$

Here λ is an overall normalization, ρ sets the scale and γ the exponent of a power-law decaying tail.

In the numerical tests shown here we took $\|\hat{H}\| = 1$ and 512 eigenvalues $\{\omega_k\}$ distributed over the whole spectrum. The two model response functions we consider are

$$\begin{aligned} S_A(\omega) &= \sum_k \frac{S_{peak}(\omega_k)}{\sum_k S_{peak}(\omega_k)} \delta(\omega - \omega_k) \\ S_B(\omega) &= \sum_k \frac{S_{peak}(\omega_k) + S_{tail}(\omega_k)}{\sum_k (S_{peak}(\omega_k) + S_{tail}(\omega_k))} \delta(\omega - \omega_k) \end{aligned} \quad (2.37)$$

where the normalization in the denominators is chosen in order to guarantee that $\mu_0 = 1$ for both response functions. The parameters for the response functions are

$$\xi = \omega_{thr} = -0.95 \quad \alpha = 5 \quad \beta = 0.05 \quad \lambda = 1 \quad \rho = 0.002 . \quad (2.38)$$

We show the two model response function in Fig. 2.1. The main panel shows the two models $S_A(\omega)$ and $S_B(\omega)$ over the entire frequency range, the presence of the tail in $S_B(\omega)$ is easily seen. In the inset we show instead the two response function in linear scale for the energy range of interest $[-1, -0.8]$, corresponding to the choice $\delta\nu = 0.2$. In this range the two responses appear almost identical. In the figure we also report the obtained values for the mean $\mu \equiv \mu_1$ and the square root of the variance σ for the two models

$$\begin{aligned} \mu^A &= -0.911 & \sigma^A &= 0.031 \\ \mu^B &= -0.907 & \sigma^B &= 0.067 \end{aligned} \quad (2.39)$$

The presence of the high-frequency tail does not modify appreciably the mean value while the variance is increased by more then a factor of two.

Due to the presence of delta function peaks, which are generated by the discreteness of the spectrum of a finite matrix, in general we cannot directly compare the integral transform Φ with the originating response function $S(\omega)$ since the normalization conventions are different. For the discrete response function we have

$$\int_{-1}^1 d\omega S(\omega) = \sum_k S(\omega_k) = 1 , \quad (2.40)$$

while for the continuous integral transform instead

$$\int_{-1}^1 d\omega \Phi(\omega) = 1 . \quad (2.41)$$

In order to directly compare the two, we approximate the integral over frequencies with the finite difference

$$\int_{-1}^1 d\omega \Phi(\omega) \approx \delta\omega \sum_{n=0}^{2/\delta\omega} \Phi(n\delta\omega - 1) , \quad (2.42)$$




Figure 2.1: *Response Functions*. The two model response functions shown as a function of frequency: black dots are for $S_A(\omega)$ while red squares for $S_B(\omega)$. The inset shows the same responses but on a narrower range and with linear scale. Also indicated are the mean μ and the square root of the variance σ for both distributions.




Figure 2.2: *Integral Transforms.* Comparison between the model response function $S_A(\omega)$ and $S_B(\omega)$ (black dots and red squares respectively) with their rescaled integral transforms $\delta\omega\Phi_A(\omega)$ and $\delta\omega\Phi_B(\omega)$ (green solid line and turquoise dotted line respectively). The inset shows the same curves but over the full energy range.

and choose $\delta\omega = 2/512$ in order to match the average spacing between the eigenvalues. In the following we will then consider the dimensionless quantity $\delta\omega\Phi$ which implies the choice $\Omega = \delta\omega$ for the energy scale of the error. We further choose the parameters of the integral kernel as

$$\Delta = 0.02 \quad \Sigma = 0.01 . \quad (2.43)$$

Fig. 2.2 shows a comparison between $S_A(\omega)$ and $S_B(\omega)$ with their rescaled integral transforms $\delta\omega\Phi_A(\omega)$ (shown as a solid green curve) and $\delta\omega\Phi_B(\omega)$ (shown as a dotted

turquoise curve). We can clearly see that the rescaling described above allows us to compare them directly.

Since the truncation error ϵ_N and the sampling error ϵ_S are more standard and the bounds described in this work are not novel, we concentrate in the following in the analysis of the systematic error ϵ_P coming from the need to perform a periodic extension of the integral kernel.

We present in Fig. 2.3 a direct comparison between the ideal (scaled) integral transform $\delta\omega\Phi_B(\omega)$ for model B (turquoise solid lines) with the approximations obtained by applying a periodic extension to the integral kernel with different choices of periods P (cf. Eq. (2.11) above). In all cases we set $\epsilon_P = 0.01$ for the target error. The approximation in the bottom panel, denoted $\delta\omega\Phi_B(\omega)$ (blue dotted line), is obtained using the conservative choice from Eq. (2.28) which is valid in general. We see clearly that the effect is to push the replicas outside the whole range of the Hamiltonian, which in our case is $[-1, 1]$, and therefore no appreciable change affects the transform in the required energy range $[\nu_{\min}, \nu_{\max}] = [-1, 0.8]$: the maximum observed deviation is $\approx 10^{-8}$. To highlight the location of this energy range, we have shaded in grey the region of energies outside of it. Also, for ease of visualization, we have presented results in a larger range of energies in order to be able to see where the replicas are located. The top panel shows instead the periodic extensions obtained using the improved bound on P obtained in this work which uses directly prior information about the first and second energy moment. It is apparent that now the integral transform outside the energy window of interest is severely distorted due to the presence of the periodic replicas. However, inside the required region, the maximum deviations are well below the value $\epsilon_P = 0.01$ chosen as target: we observe in fact an error $\approx 10^{-4}$. Very similar results have been observed for the approximations of the integral transform of model A . Before analyzing the relationship between target error and the empirically observed deviation, we want to point out the the reduction




Figure 2.3: *Periodic Extension.* Comparison between the ideal (scaled) integral transforms $\delta\omega\Phi_B(\omega)$ (turquoise solid lines) and different periodic extensions. The bottom panel shows the approximation obtained using P for a generic response denoted as $\delta\omega\Phi_B^\chi(\omega)$ (blue dotted line) while the top panel shows the approximation obtained with a smaller P found using the known energy variance σ^2 and denoted as $[\sigma]\delta\omega\Phi_B^\chi(\omega)$ (brown solid line). The shaded grey areas are energies outside the interval $[\nu_{\min}, \nu_{\max}]$.

in the value of the period are

$$\frac{P_A^\sigma}{P^{gen}} \approx 0.111 \quad \frac{P_B^\sigma}{P^{gen}} \approx 0.14 , \quad (2.44)$$

for the two model responses respectively. Here we denoted by P^{gen} the value obtained without information about the energy moments and with P^σ the value obtained using the knowledge of the energy variance. In terms of the required number of moments to guarantee a total approximation error less than $\varepsilon_P + \varepsilon_N = 0.02$ we find using Eq. (2.26)

$$N^{gen} = 218 \quad N_A^\sigma = 25 \quad N_B^\sigma = 31 . \quad (2.45)$$

In order to better understand the above observation that the measured error seems to be much smaller than the target one, which employs a possibly not tight bound on the possible error, we present in Fig. 2.4 an analysis of the scaling of the empirical error with the one set as target. The bottom panel shows how the approximation error in the scaled integral transform changes as a function of the target error for the cases where we use the improved estimate for the period P^σ which uses the energy variance. The solid black line shows the upper-bound employed to control the error while the solid green line and solid turquoise line show the observed maximum errors for both models respectively. We can see that, as expected, the real error is always below the bound we use but that the two models show a very different scaling for small target errors: the integral transform of model A has an error which decreases super-polynomially with target error while for model B the bound is saturated for small enough values of target error. We can understand this behavior in terms of the high-energy tail of the original response functions: as can be seen from the inset of Fig. 2.2 the response function for model B is dominated by the slowly decaying tail for energies outside the range of interest $[\nu_{\min}, \nu_{\max}] = [-1, -0.8]$ while the strength of model A decays much faster. Once the target error becomes comparable to the strength in the tail of model B then one is forced to consider $P^\sigma = P^{gen}$ as for a general response. On the other hand, for model A as soon as the value of P^σ exceeds the range




Figure 2.4: *Error Scaling*. The top panel shows, as a function of the target error, the fractional reduction in the period from the general case value $P = P^{gen}$ obtained using information about the energy variance: the green line shows P_A^σ/P for model A while the turquoise is P_B^σ/P for model B . The bottom panel shows the observed approximation error for the two models as a function of the target error: the green line for model A and the turquoise line for model B . The solid black line represents the best error bound we derived for this approximation (see Section 2.A.2).

of frequencies of interest than the error decays quickly following the strength response. These results show very clearly the important effect produced by slowly decaying tails in the response function and how their presence is automatically captured by the tail bounds employed in this work. The observation that the error bound is essentially saturated for small errors in model B suggests that it is unlikely to be possible to considerably improve our estimate for the optimal period P^σ in general. Finally, for not too small values of the target error $\gtrsim 10^{-3}$ our error bound overestimated the real error by more than an order of magnitude. In this regime it is possible that improved estimates using higher central moments (c.f. Eq. (2.34) and Eq. (2.93)) would be able to increase the savings in computational cost even further.

2.5 Discussion

In this work we have extended the Gaussian Integral Transform method [115] by employing a Fourier basis for the polynomial expansion of the integral transform of the spectral density. This replaces the need to estimate Chebyshev moments of the Hamiltonian with analogous Fourier moments that can be evaluated as expectation values of the real-time evolution operator $\exp(-it\hat{H})$ (c.f. Eq. (2.15)). Importantly, evaluation of both of these moments can be achieved efficiently using existing quantum algorithms. The main advantage of employing a Fourier basis is the possibility of using prior information about the spectral density, in the form of its first few energy moments, to reduce the number of observable needed for an accurate reconstruction of the spectrum. Most of our derivation and numerical examples are focused on the reasonable assumption that only the mean and variance of the energy spectrum are available to the user but we also provide improved bounds in case more moments are available.

The technique presented here can allow for orders of magnitude speed-ups in the evaluation of the spectral density. In order to give some concrete examples in nuclear physics, we take the Hamiltonian derived from Lattice EFT and used in previous cost

estimates for quantum computations of the response function (see [38]) to model a medium mass nucleus with $A = 40$. As a first example we consider a calculation of the dipole response of ^{40}Ca for excitation energies up to 100 MeV, with resolution $\Delta = 1$ MeV, small tail contributions $\Sigma = 0.01$ and approximation error $\varepsilon = 0.01$ (we take $\Omega = \Delta$ as scale). Using the experimental values from Ref. [141] (see also [142]) we estimate $\mu_1^{GT} \approx 20$ MeV and $\sigma^{GT} \approx 22$ MeV. With these values we find more than two orders of magnitude reduction in the number of moments required

$$N_{GT}^{gen} = 42372 \quad N_{GT}^{\sigma} = 339 . \quad (2.46)$$

As a second example we consider instead the simulation of longitudinal response in quasi-elastic electron-nucleus scattering at momentum transfer q . Typical values of the moments are $\mu_1^{QE} = q^2/2M$, with M the nucleon mass, and $\sigma^{QE} \approx k_F \approx 250$ MeV, with k_F the Fermi momentum [143, 144]. For a momentum transfer $q = 400$ MeV and a maximum excitation energy $\delta\nu = 400$ MeV we find

$$N_{QE}^{gen} = 42372 \quad N_{QE}^{\sigma} = 838 . \quad (2.47)$$

An expected saving by a factor ≈ 50 .

It is likely that for calculations in this regime the simple low-energy model used in Ref. [38] might not be suitable. Due to the expected increase in the Hamiltonian norm, we would expect a larger relative saving in number of moments in cases where higher resolution Hamiltonians with a larger momentum cut-off are employed.

Lastly we want to comment on the fact that the total computational cost controlled by the maximum evolution time $T = N\delta_t$ still scales as $\mathcal{O}(1/\Delta)$ and we therefore have no violation of the no-fast-forwarding theorem [26]. Interestingly for Hamiltonians that can be fast-forwarded [26, 145] the computational cost is no longer bounded by T but by N instead. For applications like the reconstruction of the spectral function or the calculation of thermodynamic observables like in Ref. [136] this will be completely given by classical cost. One can also employ the construction presented here to prepare

states in a given energy window by performing the summation in Eq. (2.12) (or more correctly it's finite N approximation) coherently on a quantum device using the Linear Combination of Unitaries strategy [146]. This is similar to the energy filter proposed in Ref. [147] or the original GIT [115] (if used coherently) and the explicit appearance of the energy variance in our cost estimates can prove useful in reducing the cost for some situations. A similar procedure could be employed as a subroutine to the Verified Phase Estimation [148] in order to estimate general expectation values. We leave a more thorough exploration of these possibilities to future work on spectral filters.

2.A Derivation of Error Bounds

In this appendix we provide a full derivation of the error bounds used in the main text. In order to simplify the notation we will denote the Gaussian kernel as

$$G_\nu(\omega) = \frac{1}{\sqrt{2\pi}\Lambda} \exp\left(-\frac{(\nu - \omega)^2}{2\Lambda^2}\right), \quad (2.48)$$

and it's periodic extension with period P as

$$G_\nu^P(\omega) = \sum_{k=-\infty}^{\infty} G_\nu(\omega + kP) = \sum_{k=-\infty}^{\infty} G_{\nu+kP}(\omega). \quad (2.49)$$

We start with the error introduced by using the periodic extension of the kernel

$$|\Phi^x(\nu) - \Phi(\nu)| = \epsilon_P(\nu). \quad (2.50)$$

2.A.1 Periodic extension in the general case

We can write this difference explicitly as

$$\begin{aligned}
\epsilon_P(\nu) &= \left| \int d\omega G_\nu^P(\omega) S(\omega) - \int d\omega G_\nu(\omega) S(\omega) \right| \\
&= \left| \int d\omega \left(\sum_{k=-\infty}^{\infty} G_{\nu+kP}(\omega) - G_\nu(\omega) \right) S(\omega) \right| \\
&= \left| \int d\omega \sum_{k \neq 0} G_{\nu+kP}(\omega) S(\omega) \right| \\
&= \int d\omega \sum_{k \neq 0} G_{\nu+kP}(\omega) S(\omega)
\end{aligned} \tag{2.51}$$

where in the last line we used the fact that both $S(\omega)$ and the Gaussian kernel are positive definite. We can now use the fact that $S(\omega) = 0$ for frequencies outside the energy spectrum so that

$$\epsilon_P(\nu) = \int_{-\|\hat{H}\|}^{\|\hat{H}\|} d\omega \sum_{k \neq 0} G_{\nu+kP}(\omega) S(\omega) . \tag{2.52}$$

At this point we can use the fact that $S(\omega)$ is integrable, and in particular we can write

$$\int_{-\|\hat{H}\|}^{\|\hat{H}\|} d\omega S(\omega) = \int d\omega S(\omega) = \mu_0 , \tag{2.53}$$

the zeroth-moment (cf. Eq. (2.9) in the main text). Using this result one can then bound the error with

$$\epsilon_P(\nu) \leq \mu_0 \sup_{\omega \in [-\|\hat{H}\|, \|\hat{H}\|]} \sum_{k \neq 0} G_{\nu+kP}(\omega) . \tag{2.54}$$

It is convenient to rewrite the summation as follows

$$[G_{\nu+P}(\omega) + G_{\nu-P}(\omega)] + \sum_{k=2}^{\infty} [G_{\nu+kP}(\omega) + G_{\nu-kP}(\omega)] ,$$

since the second term can be bounded easily using

$$\sum_{k=2}^{\infty} G_{\nu+kP}(\omega) \leq \int_1^{\infty} dx G_{\nu+XP}(\omega) , \tag{2.55}$$

$$\begin{aligned}
\sum_{k=2}^{\infty} G_{\nu+kP}(\omega) &\leq \int_1^{\infty} dx G_{\nu+XP}(\omega) \\
&= \frac{1}{\sqrt{2\pi}\Lambda} \int_1^{\infty} dx e^{-\frac{(\omega-\nu-xP)^2}{2\Lambda^2}} = \frac{1}{2P} \operatorname{erfc} \left(\frac{P-\omega+\nu}{\sqrt{2}\Lambda} \right)
\end{aligned} \tag{2.56}$$

so that, including the contribution with $-k$, we find

$$\begin{aligned}
\sum_{k=2}^{\infty} [G_{\nu+kP}(\omega) + G_{\nu-kP}(\omega)] &\leq \frac{1}{2P} \operatorname{erfc} \left(\frac{P-\omega+\nu}{\sqrt{2}\Lambda} \right) \\
&\quad + \frac{1}{2P} \operatorname{erfc} \left(\frac{P+\omega-\nu}{\sqrt{2}\Lambda} \right)
\end{aligned} \tag{2.57}$$

We can now proceed to use the bound in Eq. (2.17) of the main text to write

$$\epsilon_P(\nu) \leq C \sup_{\omega \in [-\|\hat{H}\|, \|\hat{H}\|]} (G_{\nu+P}(\omega) + G_{\nu-P}(\omega)) , \tag{2.58}$$

where the constant factor C is given by

$$C = \mu_0 \left(1 + \sqrt{\frac{\pi}{2}} \frac{\Lambda}{P} \right) . \tag{2.59}$$

In order to turn this into a useful bound, we will consider the largest error that can occur for any

$$\nu \in [\nu_{\min}, \nu_{\max}] \subseteq [-\|\hat{H}\|, \|\hat{H}\|] , \tag{2.60}$$

which directly implies that

$$\omega - \nu \in [-\|\hat{H}\| - \nu_{\max}, \|\hat{H}\| - \nu_{\min}] \subseteq [-2\|\hat{H}\|, 2\|\hat{H}\|] . \tag{2.61}$$

First, notice that one of two Gaussians in Eq. (2.58) will always dominate over the other. To see this we will consider two cases separately: first let's take

$$\omega - \nu \in [-\|\hat{H}\| - \nu_{\max}, 0] . \tag{2.62}$$

In this range of values we have

$$\begin{aligned}
|P - \omega + \nu| &= P - \omega + \nu \geq P > 0 \\
|P + \omega - \nu| &= |P - |\omega - \nu|| \begin{cases} > 0 & \text{for } P > \|\hat{H}\| + \nu_{\max} \\ \geq 0 & \text{for } P \leq \|\hat{H}\| + \nu_{\max} \end{cases}
\end{aligned} \tag{2.63}$$

For this range, the second Gaussian centered in $\nu - P$ dominates and we need to have $P > \|\hat{H}\| + \nu_{\max}$ in order to prevent the exponent to go to zero. Therefore

$$\begin{aligned} G_{\nu+P}(\omega) + G_{\nu-P}(\omega) &\leq 2G_{\nu-P}(\omega) \\ &\leq \frac{2}{\sqrt{2\pi}\Lambda} \exp\left(-\frac{(P - \|H\| - \nu_{\max})^2}{2\Lambda^2}\right). \end{aligned} \quad (2.64)$$

If we take the complementary range of values

$$\omega - \nu \in [0, \|\hat{H}\| - \nu_{\min}] , \quad (2.65)$$

we find instead the the first Gaussian dominate and, for $P > \|\hat{H}\| - \nu_{\min}$ we find the useful bound

$$\begin{aligned} G_{\nu+P}(\omega) + G_{\nu-P}(\omega) &\leq 2G_{\nu+P}(\omega) \\ &\leq \frac{2}{\sqrt{2\pi}\Lambda} \exp\left(-\frac{(P - \|H\| + \nu_{\min})^2}{2\Lambda^2}\right). \end{aligned} \quad (2.66)$$

These results suggest that we should take

$$P = (1 + \eta)\|\hat{H}\| + \max[|\nu_{\min}|, \nu_{\max}] , \quad (2.67)$$

for some appropriate $\eta > 0$. For this choice we have in fact

$$\epsilon_P(\nu) \leq \frac{2\mu_0}{\sqrt{2\pi}\Lambda} \left(1 + \sqrt{\frac{\pi}{2}} \frac{\Lambda}{P}\right) e^{-\frac{\eta^2 \|\hat{H}\|^2}{2\Lambda^2}}. \quad (2.68)$$

In order to simplify the calculation of a good value for η that would guarantee $\epsilon_P(\nu)\Omega < \varepsilon_P$, for some (dimensionless) target error tolerance $\varepsilon_P > 0$, we use the simple lower bound $P > \|\hat{H}\|$ and find

$$\eta \geq \frac{\sqrt{2}\Lambda}{\|\hat{H}\|} \sqrt{\log\left(\sqrt{\frac{2}{\pi}} \frac{\mu_0}{\varepsilon_P} \frac{\Omega}{\Lambda} \left(1 + \sqrt{\frac{\pi}{2}} \frac{\Lambda}{\|\hat{H}\|}\right)\right)}. \quad (2.69)$$

Neglecting sub-leading logarithmic factors we find therefore the asymptotic scaling

$$\chi = \frac{P}{\|\hat{H}\|} = \tilde{\mathcal{O}}\left(1 + \frac{\Lambda}{\|\hat{H}\|} \sqrt{\log\left(\frac{\mu_0\Omega}{\varepsilon_P\Lambda}\right)}\right). \quad (2.70)$$

A more practical bound which uses the conventions of the main text is the following one

$$\chi = 2 + \frac{\sqrt{2}\Lambda}{\|\hat{H}\|} \sqrt{\log \left(\frac{2\Omega}{\varepsilon_P \Lambda} \right)}, \quad (2.71)$$

which is the result quoted in the main text.

2.A.2 Periodic extension with moment information

We now turn to show how to obtain an improved error bound using information about the energy moments. As before, and without loss of generality, we will consider $\nu \in [\nu_{\min}, \nu_{\max}]$. For the following manipulations, it will be convenient to change variables and define

$$\nu_{\min} = \mu_1 - \Omega_{\min} \quad \nu_{\max} = \mu_1 + \Omega_{\max}, \quad (2.72)$$

with μ_1 the first energy moment which gives us information about the average energy in the spectral function and $\Omega_{\min}, \Omega_{\max}$ both positive. At this point we start from the expression for the error $\epsilon_P(\nu)$ obtained in Eq. (2.51) before and rewrite it as follows

$$\begin{aligned} \epsilon_P(\nu) &= \int d\omega \sum_{k \neq 0} G_{\nu+kP}(\omega) S(\omega) \\ &= \int_{-\infty}^{\mu_1 - \alpha\sigma} d\omega \sum_{k \neq 0} G_{\nu+kP}(\omega) S(\omega) \\ &\quad + \int_{\mu_1 - \alpha\sigma}^{\mu_1 + \alpha\sigma} d\omega \sum_{k \neq 0} G_{\nu+kP}(\omega) S(\omega) \\ &\quad + \int_{\mu_1 + \alpha\sigma}^{\infty} d\omega \sum_{k \neq 0} G_{\nu+kP}(\omega) S(\omega) \\ &= \epsilon_1(\nu) + \epsilon_2(\nu) + \epsilon_3(\nu), \end{aligned} \quad (2.73)$$

for some $\alpha > 0$ and where we have denoted by σ the square root of the variance

$$\sigma = \sqrt{\mu_2 - \mu_1^2}. \quad (2.74)$$

The central integral can be bounded in the same way we obtained the result in the previous section, the main difference is that now we take the interval

$$\omega - \nu \in [-\alpha\sigma - \Omega_{\max}, \alpha\sigma + \Omega_{\min}] . \quad (2.75)$$

In this range we can bound the central contribution with

$$\epsilon_2(\nu) \leq \frac{2\mu_0}{\sqrt{2\pi}\Lambda} \left(1 + \sqrt{\frac{\pi}{2}} \frac{\Lambda}{P} \right) e^{-\frac{\eta^2 \alpha^2 \sigma^2}{2\Lambda^2}} , \quad (2.76)$$

where now we took the period to be

$$P = (1 + \eta)\alpha\sigma + \max[\Omega_{\min}, \Omega_{\max}] , \quad (2.77)$$

for some $\eta > 0$. Note that, in order for the exponential to become small, we would need $P > \eta\alpha\sigma > \Lambda$ so that we can use the simpler bound

$$\epsilon_2(\nu) \leq \frac{2\mu_0}{\sqrt{2\pi}\Lambda} \left(1 + \sqrt{\frac{\pi}{2}} \right) e^{-\frac{\eta^2 \alpha^2 \sigma^2}{2\Lambda^2}} , \quad (2.78)$$

which will incur at most a logarithmic cost. In order to guarantee this term to be smaller than $\varepsilon_P/(2\Omega)$, for dimensionless $\varepsilon_P > 0$, we can take (cf. Eq. (2.69))

$$\eta = \frac{\sqrt{2}\Lambda}{\alpha\sigma} \sqrt{\log \left(\sqrt{\frac{8}{\pi}} \frac{\mu_0}{\varepsilon_P} \frac{\Omega}{\Lambda} \left(1 + \sqrt{\frac{\pi}{2}} \right) \right)} . \quad (2.79)$$

In order to bound the other two terms, we first bound the sum over k with an integral using

$$\sum_{k=1}^{\infty} G_{\nu+kP}(\omega) \leq \int_0^{\infty} dx G_{\nu+XP}(\omega) , \quad (2.80)$$

together with a similar one for the negative terms

$$\begin{aligned} \sum_{k=1}^{\infty} G_{\nu-kP}(\omega) &\leq \int_0^{\infty} dx G_{\nu-XP}(\omega) \\ &= \int_{-\infty}^0 dx G_{\nu+XP}(\omega) , \end{aligned} \quad (2.81)$$

so that summing them together we find the bound

$$\sum_{k \neq 0}^{\infty} G_{\nu+kP}(\omega) \leq \int_{-\infty}^{\infty} dx G_{\nu+xP}(\omega) = \frac{1}{P}. \quad (2.82)$$

The sum of the two missing terms $\epsilon_{13}(\nu) = \epsilon_1(\nu) + \epsilon_3(\nu)$ is thus bounded by the following

$$\begin{aligned} \epsilon_{13}(\nu) &\leq \frac{1}{P} \left(\int_{-\infty}^{\mu-\alpha\sigma} d\omega S(\omega) + \int_{\mu+\alpha\sigma}^{\infty} d\omega S(\omega) \right) \\ &= \frac{1}{P} \left(1 - \int_{\mu-\alpha\sigma}^{\mu+\alpha\sigma} d\omega S(\omega) \right) \\ &\leq \frac{\sigma^2}{P\alpha^2\sigma^2} = \frac{1}{P\alpha^2} \\ &\leq \frac{1}{(1+\eta)\alpha^3\sigma}. \end{aligned} \quad (2.83)$$

Here we used the normalization of $S(\omega)$ to get to the second line, the Chebyshev bound Eq. (2.31) in the main text to get to the third and Eq. (2.77) for the last one. At this point we can in principle use the expression for η in Eq. (2.79) to find an appropriate value for α so that $\epsilon_{13}(\nu) < \varepsilon_P/(2\Omega)$. In order to better see the general trend, we instead ensure that α satisfies

$$\alpha \geq \frac{\sqrt{2}\Lambda}{\sigma} \sqrt{\log \left(3.6 \frac{\mu_0\Omega}{\varepsilon_P\Lambda} \right)}, \quad (2.84)$$

so that $\eta < 1$. We can then solve

$$\frac{1}{(1+\eta)\alpha^3\sigma} \leq \frac{1}{\alpha^3\sigma} \leq \frac{\varepsilon_P}{2\Omega}. \quad (2.85)$$

This results in the following bound

$$\alpha \geq \frac{\sqrt{2}\Lambda}{\sigma} \max \left[\sqrt{\log \left(3.6 \frac{\mu_0\Omega}{\varepsilon_P\Lambda} \right)}, \frac{\sigma^{2/3}\Omega^{1/3}}{\Lambda} \frac{0.9}{\varepsilon_P^{1/3}} \right]. \quad (2.86)$$

This result can be inserted directly into Eq. (2.77) to obtain the final estimate for the period P required.

We now use the conventions described in the main text to obtain a more intuitive result and distinguish explicitly two different regimes. For $2\sigma \geq \Lambda$ the second term will dominate for all $\varepsilon_P \lesssim 0.6$ and for all values if we increase the numerical factor to 1.9. In this case we can take

$$\chi = \frac{2.7}{\varepsilon_P^{1/3}} \frac{\Omega^{1/3} \sigma^{2/3}}{\|\hat{H}\|} + \frac{\delta\nu}{\|\hat{H}\|} , \quad (2.87)$$

as quoted in the main text. Here we used the fact that

$$\max[\Omega_{\min}, \Omega_{\max}] \leq \Omega_{\min} + \Omega_{\max} = \nu_{\max} - \nu_{\min} = \delta\nu . \quad (2.88)$$

For large values of $\Lambda > 2\sigma$ we would need to use the full expression in Eq. (2.86) above instead.

Finally, for cases where we know the value of central moments

$$\tilde{\mu}_n = \int d\omega |\omega - \mu_1|^n S(\omega) , \quad (2.89)$$

for n higher than 2 are known, then one can use the following generalization of the Chebyshev inequality

$$Prob[|\omega - \mu_1| \geq \Gamma] \leq \frac{\widetilde{\mu}_n}{\Gamma^n} . \quad (2.90)$$

Using this additional information we can take instead

$$P = (1 + \eta) \alpha \widetilde{\mu}_n^{1/n} + \max[\Omega_{\min}, \Omega_{\max}] , \quad (2.91)$$

and can therefor choose α as follows

$$\alpha \geq \frac{\sqrt{2}\Lambda}{\widetilde{\mu}_n^{1/n}} \max \left[\sqrt{\log \left(4 \frac{\mu_0 \Omega}{\varepsilon_P \Lambda} \right)}, \left(\frac{\widetilde{\mu}_n \Omega}{\varepsilon_P} \right)^{\frac{1}{n+1}} \frac{0.9}{\Lambda} \right] . \quad (2.92)$$

For $\widetilde{\mu}_n^{1/n} \geq \Lambda$ we have the following simpler result

$$\chi = \frac{2.7}{\varepsilon_P^{1/(n+1)}} \frac{(\Omega \widetilde{\mu}_n)^{1/(n+1)}}{\|\hat{H}\|} + \frac{\delta\nu}{\|\hat{H}\|} , \quad (2.93)$$

valid up to $n = 15$. This can be advantageous for small errors provided the central moments do not grow too much.

2.A.3 Truncation error

Here we provide the details underlying the bound for the truncation error

$$\epsilon_N(\nu) = |\Phi_N^x(\nu) - \Phi^x(\nu)|, \quad (2.94)$$

used in the main text. In order to simplify the notation we will use m_n^x to denote the moments (see Eq. (2.7)). For our Fourier polynomials these are

$$\begin{aligned} m_n^x &= \int d\omega S(\omega) \exp\left(-i \frac{2\pi}{\chi \|\hat{H}\|} n\omega\right) \\ &= \langle \Psi_0 | \hat{O}^\dagger \exp\left(-i \frac{2\pi}{\chi \|\hat{H}\|} n\hat{H}\right) \hat{O} | \Psi_0 \rangle, \end{aligned} \quad (2.95)$$

In addition, note that $|m_n^x| \leq \mu_0$ with μ_0 the zeroth energy moment. This can be easily seen by normalizing the state $\hat{O}|\Psi_0\rangle$ with $\sqrt{\mu_0}$ and using the fact that the evolution operator is unitary. We can now express explicitly the truncation as follows

$$\epsilon_N(\nu) = \frac{1}{\chi \|\hat{H}\|} \left| \sum_{n=N+1}^{\infty} (g_n^x(\nu) m_n^x + g_{-n}^x(\nu) m_{-n}^x) \right|, \quad (2.96)$$

with $g_n^x(\nu)$ the Fourier coefficients from Eq. (2.20). Using the bound on the moments m_n^x we then find immediately

$$\begin{aligned} \epsilon_N(\nu) &\leq \frac{\mu_0}{\chi \|\hat{H}\|} \sum_{n=N+1}^{\infty} (|g_n^x(\nu)| + |g_{-n}^x(\nu)|) \\ &= \frac{2\mu_0}{\chi \|\hat{H}\|} \sum_{n=N+1}^{\infty} \exp\left(-\frac{2\pi^2 \Lambda^2}{\chi^2 \|\hat{H}\|^2} n^2\right) \\ &\leq \frac{2\mu_0}{\chi \|\hat{H}\|} \int_N^{\infty} dx \exp\left(-\frac{2\pi^2 \Lambda^2}{\chi^2 \|\hat{H}\|^2} x^2\right) \\ &= \frac{\mu_0}{\sqrt{2\pi} \Lambda} \operatorname{erfc}\left(N \frac{\sqrt{2\pi} \Lambda}{\chi \|\hat{H}\|}\right) \\ &\leq \frac{\mu_0}{\sqrt{2\pi} \Lambda} \exp\left(-N^2 \frac{2\pi^2 \Lambda^2}{\chi^2 \|\hat{H}\|^2}\right) \end{aligned} \quad (2.97)$$

where we used the triangle inequality on the first line, the bound on a sum with an integral in the third and Eq. (2.17) in the last. Note that the bound does not depend on the frequency ν anymore. We can now find the value for N that would guarantee $\epsilon_N(\nu)\Omega \leq \varepsilon_N$ for some (dimensionless) target error tolerance $\varepsilon_N > 0$. The result is

$$N \geq \frac{\chi \|\hat{H}\|}{\sqrt{2\pi}\Lambda} \sqrt{\log \left(\frac{\mu_0 \Omega}{\sqrt{2\pi}\Lambda \varepsilon_N} \right)}. \quad (2.98)$$

This is the result used in the main text.

2.A.4 Statistical error

We now turn to the discussion of the bound on the statistical error in the evaluation of the Fourier moments m_n^χ . For the moment we will assume we have estimated each moment with a fixed error ε_M which for simplicity we take to be equal for all moments (thanks to the rapid decrease in the coefficients an adaptive strategy be more advantageous). We also neglect the error on the zeroth moment m_0^χ since its value is known beforehand. Since errors on different moments are independent, we add the error contributions in quadrature to find

$$\begin{aligned} \epsilon_S^2 &\leq \frac{\varepsilon_M^2 \mu_0^2}{\chi^2 \|\hat{H}\|^2} \sum_{n=1}^N \left(|g_n^\chi(\nu)|^2 + |g_{-n}^\chi(\nu)|^2 \right) \\ &= 2 \frac{\varepsilon_M^2 \mu_0^2}{\chi^2 \|\hat{H}\|^2} \sum_{n=1}^N \exp \left(-\frac{4\pi^2 \Lambda^2}{\chi^2 \|\hat{H}\|^2} n^2 \right) \\ &\leq 2 \frac{\varepsilon_M^2 \mu_0^2}{\chi^2 \|\hat{H}\|^2} \int_0^N dx \exp \left(-\frac{4\pi^2 \Lambda^2}{\chi^2 \|\hat{H}\|^2} x^2 \right) \\ &< 2 \frac{\varepsilon_M^2 \mu_0^2}{\chi^2 \|\hat{H}\|^2} \int_0^\infty dx \exp \left(-\frac{4\pi^2 \Lambda^2}{\chi^2 \|\hat{H}\|^2} x^2 \right) \\ &= \frac{\varepsilon_M^2 \mu_0^2}{2\pi \chi \|\hat{H}\| \Lambda}, \end{aligned} \quad (2.99)$$

where we used the fact that the variance of the moments is less than μ_0^2 and the same procedure employed in Eq. (2.97). In order to attain a total expected error $\epsilon_S \leq \varepsilon_S/\Omega$

we then need to take

$$\varepsilon_M < \frac{\varepsilon_S}{\mu_0} \frac{\sqrt{2\pi\chi\|\hat{H}\|\Lambda}}{\Omega} \quad (2.100)$$

resulting in an expected number of samples

$$S > 2N \frac{\Omega^2 \mu_0^2}{2\pi\chi\|\hat{H}\|\Lambda\varepsilon_S^2} . \quad (2.101)$$

The additional factor of 2 in the numerator comes from the need to evaluate separately the real and imaginary part of each moment separately. If we want to ensure this is sufficient with high probability we can use Markov's to ensure the probability that the error is below ε_S^2 is larger than $2/3$ by increasing the target error ε_S by a factor of at least $\sqrt{3}$ and then use the Chernoff bound and majority voting to increase the probability to $1 - \delta$ with logarithmic effort. For instance

$$S = N \frac{\Omega^2 \mu_0^2}{\chi\|\hat{H}\|\Lambda\varepsilon_S^2} \log\left(\frac{2}{\delta}\right) , \quad (2.102)$$

will be enough for a confidence level $1 - \delta$. Together with the bound from Eq. (2.98) this shows that the number of samples is independent from the number of terms N .

The treatment above assumes errors are completely uncorrelated which might not necessarily be the case due to the need of controlling systematic errors with e.g. error mitigation techniques. For a more conservative error estimate we consider instead a bound to ε_S obtained by summing the individual errors in absolute value. Following the same procedure used above we find the final result

$$S = N \frac{\Omega^2 \mu_0^2}{\Lambda^2 \varepsilon_S^2} \log\left(\frac{2}{\delta}\right) , \quad (2.103)$$

quoted in the main text. We want to conclude this appendix with a similar result regarding the original Chebyshev based GIT from Ref. [115]. The estimate of the sample complexity reported there didn't use the strategy employed here and as a result the original work gave an estimate $S = \mathcal{O}(N^3)$ which was somewhat pessimistic. Using in fact the present strategy, together with the results in App.D.3 of [115] and

restoring the energy dimensions in order to be compatible with our current conventions we can show that a number of samples given by

$$S_{Cheb} = 2N \frac{\Omega^2}{\Lambda^2 \varepsilon_S^2} \log \left(\frac{2}{\delta} \right) , \quad (2.104)$$

are enough to control the statistical errors.

Chapter 3

STABILIZER SCARS

This work has been published as “Stabilizer Scars” by Jeremy Hartse, Luasz Fidkowski and Niklas Mueller, in Physical Review Letters [149]

3.1 Introduction

Unprecedented advances in controlling isolated quantum systems [150, 21, 22] have enabled the study of non-equilibrium phenomena in complex many-body systems, including thermalization, a pivotal subject in many fields [151, 51, 152, 153]. As discussed in 1.5, according to the eigenstate thermalization hypothesis (ETH) [52, 53], non-integrable isolated systems thermalize, with quantum ergodicity ensuring that states explore a vast Hilbert space regardless of the initial state.

However, recent experiments have identified so-called *quantum many-body scars* (QMBS), where quantum ergodicity and the ETH are violated despite the system being strongly coupled and non-integrable [54, 55, 56, 57, 58]. This has sparked significant theoretical and experimental interest [78, 60, 77, 76, 73, 79, 72, 154, 59, 71, 70, 69, 68, 67, 88, 62, 74, 75, 61, 80, 63, 64, 65, 91, 81, 155, 85, 84, 86, 82], but many questions are still open: For instance, the analytic mechanisms behind QMBS, their stability in the thermodynamic limit and their fate as generic phenomena, beyond finely-engineered synthetic quantum systems, or use in quantum information and computing, are presently unclear.

Seeking insight into the mechanisms behind QMBS, in this chapter we investigate a quasi-1d limit of a strongly coupled and non-integrable $2 + 1$ d model of a lattice gauge theory (LGT). LGTs are an important target for quantum simulators and

computers because of their relevance in high-energy and nuclear physics [14, 156, 12, 13, 157], where thermalization is a central objective [158, 159, 160, 161], for topological phases [162, 163], or quantum error correction [164, 165, 166].

Our study, based on an Ising-LGT duality and fermionization, uncovers exact analytic scar solutions for arbitrary couplings and in the thermodynamic limit. Importantly, the scar subspace of the model is spanned by certain *stabilizer states* [93, 95] with zero magic resource [96, 97], revealing a connection between complexity and quantum ergodicity.

3.2 Model

We consider $\mathbb{Z}_2$ LGT [167, 168] with Hamiltonian,

$$H = - \sum_p \prod_{\nu \in p} \sigma_\nu^x - g \sum_\nu \sigma_\nu^z \quad (3.1)$$

where $\sigma^{x/z}$ are Pauli matrices, ν are the links of a square lattice in $2 + 1$ dimensions with periodic boundary conditions, and p are plaquettes, see Fig. 3.1. Local Gauss law operators, $G_s \equiv \prod_{\nu \in s} \sigma_\nu^z$, where $\nu \in s$ are links ν originating from a lattice site s , see Fig. 3.1, and two ribbon operators $V_{x,y} \equiv \prod_{\nu \in \mathcal{L}_{x/y}} \sigma_\nu^z$, where $\mathcal{L}_{x/y}$ wind periodically around x/y , define superselection sectors, $[H, G_s] = [H, V_x] = [H, V_y] = 0$ [169]. We consider a lattice of $L \times 2$ plaquettes and, w.l.o.g., set $V_x = V_y = 1$, and consider the combined $+1$ eigenspace of all G_s ; see more in Appendix 3.B.

The model described by Eq. (3.1) is dual to a $(2 + 1)$ d transverse-field Ising model with periodic boundary conditions when restricted to the parity-even subspace, $\prod_p X_p = 1$, via the map $X_p \equiv \prod_{\nu \in p} \sigma_\nu^x$ and $Z_p Z_{p+\hat{a}} = \sigma_\nu^z$, where p and $p + \hat{a}$ label the plaquettes adjacent to the link ν ,

$$H^{\text{dual}} = - \sum_{p=1}^L [X_p + X_{p+L}] - \tilde{g} \sum_{p=1}^L Z_p Z_{p+L} - \sum_{p=1}^L g_p [Z_p Z_{p+1} + Z_{p+L} Z_{p+L+1}], \quad (3.2)$$

where $\tilde{g} \equiv g(1 + V_y)$, and $g_p \equiv g$ for $p < L$ and $g_p = V_x g$ for $p = L$; X_p and Z_p are Pauli matrices centered on plaquettes, see Fig. 3.2(a). Our notation is such that




Figure 3.1: $\mathbb{Z}_2$ LGT. Degrees of freedom of $\mathbb{Z}_2$ LGT on the edges of a (periodic) two-dimensional square lattice; plaquette, electric field and Gauss law operators are shown.

plaquettes are numbered 1 to L from left to right in the top row and from $L + 1$ to $2L$ from left to right in the bottom row. For more details about the duality see Appendix 3.A.

3.3 Quantum Many-Body Scars

For an arbitrary value of the coupling g , and $L \times 2$ geometry, the model's eigenspace possesses a subspace of exact QMBS. For even L , this subspace is two-dimensional and spanned by the states,

$$\rho_{1/2} \equiv \prod_{p=1}^L \left[\frac{1 \pm (-1)^p Z_p Z_{p+L}}{2} \right] \left[\frac{1 - X_p X_{p+L}}{2} \right] \quad (3.3)$$

where $\rho_{1/2} \equiv |\Psi_{1/2}\rangle\langle\Psi_{1/2}|$ are (eigen-)states with exactly zero energy, $H|\Psi_{1/2}\rangle = 0$. In Appendix 3.C, we derive Eq. (3.3).

Importantly, Eq. (3.3) are *stabilizer states*, specified by the eigenvalues of $X_p X_{p+L}$ and $Z_p Z_{p+L}$ and are depicted in Fig. 3.2(b). The corresponding scar states in LGT




Figure 3.2: *Quantum Many-Body Scars in $\mathbb{Z}_2$ LGT*. (a) Illustration of Dual Ising model with stabilizers $X_p X_{p+L}$ (eigenvalues in light and dark blue) and $Z_p Z_{p+L}$ (eigenvalues in light and dark red). (b) QMBS scar solutions of a $L \times 2$ lattice where L is even, shown in terms of $X_p X_{p+L}$ and $Z_p Z_{p+L}$ eigenvalues. (c) The scar subspace for odd L is spanned by the states shown.

formulation are obtained upon replacing $X_p X_{p+L} = \prod_{\nu \in p} \sigma_\nu^x \sigma_{\nu+L}^x$ and $Z_p Z_{p+L} = \sigma_\nu^z$; additionally, on this side of the duality, Gauss's laws (and V_x, V_y) are stabilizers, i.e.,

$$\rho_{1/2} \rightarrow \rho_{1/2} \prod_s \left[\frac{1 + G_s}{2} \right] \frac{1 + V_x}{2} \frac{1 + V_y}{2}, \quad (3.4)$$

where the l.h.s. is the Ising and the r.h.s. the LGT expression. For brevity, we also use $X_p X_{p+L}$ and $Z_p Z_{p+L}$ when discussing the LGT, referring to $\prod_{\nu \in p} \sigma_\nu^x \sigma_{\nu+L}^x$ and σ_ν^z , respectively.

For odd L , the scar subspace is spanned by $4L$ stabilizer states that are depicted in Fig. 3.2(c), where $\varphi_\alpha \equiv |\varphi_\alpha\rangle\langle\varphi_\alpha|$,

$$\begin{aligned} \varphi_\alpha(k) &\equiv \prod_{q < k} \left[\frac{1 + s_\alpha (-1)^{k-q} Z_q Z_{q+L}}{2} \right] \frac{1 + t_\alpha Z_k Z_{k+L}}{2} \\ &\times \prod_{q > k} \left[\frac{1 - s_\alpha (-1)^{k-q} Z_q Z_{q+L}}{2} \right] \Phi(k). \end{aligned} \quad (3.5)$$

Here, $\alpha = 1, 2, 3, 4$ and $k = 1, \dots, L$; furthermore we abbreviate $s_1 = s_4 \equiv 1$, $s_2 = s_3 \equiv -1$, as well as $t_2 = t_4 \equiv 1$ and $t_1 = t_3 \equiv -1$. Additionally,

$$\Phi(k) \equiv \prod_{q \neq k} \left[\frac{1 - X_q X_{q+L}}{2} \right] \frac{1 + X_k X_{k+L}}{2}. \quad (3.6)$$

Unlike for even L , Eq. (3.5) are not eigenstates, but rather they define a conserved subspace in which the action of H maps scar states onto other scar states exactly while having zero overlap with non-scar states; we demonstrate this property in Appendix 3.C. The states in the LGT formulation are defined analogously to Eq. (3.4), with Gauss laws and V_x, V_y additional stabilizers. In the *End Matter* section we discuss the commutant algebra that generates the scar subspace.

3.4 Entanglement and Magic Resource

The scar basis states, Eq. (3.3) and Eq. (3.5), have zero magic resource because they are stabilizer states. We investigate their entanglement, focusing on the von Neumann entanglement entropy of a bipartition along the long (L) lattice direction,

$\rho_A = \text{Tr}_{\bar{A}}(\rho)$, where A is the subsystem and $\bar{A}$ its complement. The reduced density matrix $\rho_A \equiv \bigoplus_s \rho_A^{(s)}$ exhibits a block-diagonal structure, whose origin we explain below. As a result, the von Neumann entanglement entropy can be decomposed into distillable and symmetry contributions,

$$S_{\text{vN}} = S_{\text{vN}}^{\text{dist}} - \sum_s p_s \log(p_s), \quad (3.7)$$

where the distillable von Neumann entropy is

$$S_{\text{vN}}^{\text{dist}} = \sum_s p_s S_{\text{vN}}^{(s)}, \quad S_{\text{vN}}^{(s)} = -\text{Tr}(\tilde{\rho}_A^{(s)} \log(\tilde{\rho}_A^{(s)})), \quad (3.8)$$

with $\tilde{\rho}_A^{(s)} \equiv \rho_A^{(s)} / p_s$ and $p_s = \text{Tr}(\rho_A^{(s)})$. Although the LGT and Ising models are dual to each other, the symmetries of ρ_A and its entanglement structure generally differ between the two formulations.

In the top row of Fig. 3.3 we show the distillable entanglement for even $L = 8$ and $g = 0.9$ in both the LGT (left panel) and the Ising model (right panel). For both zero-energy scar eigenstates, the distillable entanglement vanishes. In contrast, the bottom row illustrates the same quantity for *odd* $L = 7$. Here, the eigenstates are linear superpositions of the stabilizer scar solutions, Eq. (3.5). In the LGT, the distillable entanglement remains zero for any coupling, whereas it is non-zero in the Ising dual.

3.5 Zero Distillable Entanglement in Lattice Gauge Theory Scar States

The conserved scar subspace exists in both the Ising model and the LGT, but the scars' distillable entanglement is *exactly zero for any value of coupling* only in the LGT owing to the special superselection sector structure in the latter. To explain this structure, we consider the entanglement cut shown in Fig. 3.4(a). Gauss' law implies that $\sigma_\nu^z = \pm 1$ on each boundary link $\nu \in \{\nu_1^a, \nu_2^a, \nu_1^b, \nu_2^b\}$ is a distinct superselection operator. However, the constraint $\prod_{\nu \sim \mathcal{L}_y} \sigma_\nu^z = V_y = 1$, where $\mathcal{L}_y$ is any non-trivial loop in the y direction and $\nu \sim \mathcal{L}_y$ are the links cut by $\mathcal{L}_y$, gives $\sigma_{\nu_1^{a/b}}^z \sigma_{\nu_2^{a/b}}^z = 1$.




Figure 3.3: *Quantum Many-Body Scars: Lattice Gauge Theory vs. Ising Dual*. Top row: Distillable entanglement of all eigenstates for even L , for $L = 8$ and $g = 0.9$. The scar states, given by Eq. (3.3) and shown as red stars, appear in both the LGT (left panel) and the dual Ising model (right panel) and are characterized by exactly zero distillable entanglement. Bottom row: For odd L ($L = 7$ and $g = 0.9$), the scar eigenstates form linear superpositions of the stabilizer states, Eq. (3.4). The scars exhibit zero distillable entanglement in the LGT, but not in the dual Ising model.

Consequently, there is only one superselection operator $\sigma_{\nu_1^{a/b}}^z$ associated with each boundary, which in the Ising dual equal $Z_a Z_{a+L}$ and $Z_b Z_{b+L}$. A third superselection operator arises from the constraint $\prod_{\nu \sim \mathcal{L}_x} \sigma_\nu^z = V_x = 1$. Its restriction to A , shown in Fig. 3.4(b), $\prod_{\nu \sim \mathcal{L}_x^A} \sigma_\nu^z \sim Z_a Z_b$, where $\mathcal{L}_x^A = \mathcal{L}_x \cap A$, gives the third superselection operator. We refer to Appendix 3.B for more details.

These three superselection operators result in zero distillable entanglement for the scar energy eigenstates in the LGT (but not for the Ising dual): Using the notation of Fig. 3.2(c), a scar energy eigenstate is a superposition

$$|\psi\rangle = \sum_{1 \leq k \leq L, 1 \leq i \leq 4} c_{k,i} |\phi_i(k)\rangle \quad (3.9)$$

where $\sum_{k,i} |c_{k,i}|^2 = 1$ and $|\phi_i(k)\rangle$ is a LGT scar basis state via our dictionary. We now group the terms in this superposition according to their boundary superselection sectors

$$|\psi\rangle = \sum_{\alpha, \beta = \pm 1} |\psi_{\alpha, \beta}\rangle, \quad (3.10)$$

i.e., $Z_a Z_{a+L} |\psi_{\alpha, \beta}\rangle = \alpha |\psi_{\alpha, \beta}\rangle$ and $Z_b Z_{b+L} |\psi_{\alpha, \beta}\rangle = \beta |\psi_{\alpha, \beta}\rangle$. The key observation is that, when considering a reduced density matrix on A , the four terms fall into separate superselection sectors.

Since the argument is nearly identical for all four states $|\psi_{\alpha, \beta}\rangle$, we will demonstrate it for $\alpha = 1$ and $\beta = -1$ and a and b are even. We consider

$$\begin{aligned} |\psi_{+1, -1}\rangle &= \sum_{(a < k < b) \text{ and } (k \text{ even}), i=1,4} c_{k,i} |\phi_i(k)\rangle \\ &+ \sum_{(a < k < b) \text{ and } (k \text{ odd}), i=2,3} c_{k,i} |\phi_i(k)\rangle \\ &+ c_{a,4} |\phi_4(a)\rangle + c_{b,1} |\phi_1(b)\rangle \end{aligned} \quad (3.11)$$

which represents a state where the $Z_k Z_{k+L}$ defect, as illustrated in Fig. 3.2(c), is trapped within region A – the set-up is shown in Fig. 3.4(c). Any state described by




Figure 3.4: *Quantum Many-Body Scars: Lattice Gauge Theory vs. Ising Dual II.* (a) Superselection sectors, $\sigma_\nu^z = \pm 1$ on each boundary link $\nu \in \{\nu_1^a, \nu_2^a, \nu_1^b, \nu_2^b\}$, are associated with Gauss laws. (b) An additional superselection sector, $\prod_{\nu \sim \mathcal{L}_x^A} \sigma_\nu^z \sim Z_a Z_b$, where $\mathcal{L}_x^A = \mathcal{L}_x \cap A$, arises from restricting $\prod_{\nu \sim \mathcal{L}_x} \sigma_\nu^z = V_x$ to A (similar for $\bar{A}$). (c) The $Z_p Z_{p+L}$ eigenvalue configuration, for $p \leq a$ and $p \geq b$, is shown for the state $|\psi_{+1,-1}\rangle$ from Eq. (3.11), with a defect trapped inside A . Along with the $X_p X_{p+L}$ and Gauss law constraints, the reduced density matrix of $|\psi_{+1,-1}\rangle$ on $\bar{A}$ (and A) has at most rank 2. The light red and dark red colors indicate the $Z_p Z_{p+L}$ eigenvalue as in fig 3.2. (d) While $Z_p Z_{p+L}$ constrains horizontal links in the LGT, Gauss law eigenstates require vertically separated vertical links to be in either a $+1, -1$ or $-1, +1$ configuration, providing an effective spin-1/2 degree of freedom. Imposing the $X_p X_{p+L}$ constraint leads to two states—one even, one odd under $\prod_{\nu \sim \mathcal{L}_x^A} \sigma_\nu^z$ —resembling the degenerate ground state of a 1D Ising model without a transverse field.

Eq. (3.11) shares the same stabilizer values for $X_p X_{p+L}$ and $Z_p Z_{p+L}$ for all $p \leq a$ and $p \geq b$ ¹, which imposes $2(L - b + a - 1)$ constraints. Additionally, enforcing Gauss's law in the complement region introduces $2(L - b + 1) - 1$ further constraints. For given α and β , there are $4(L + a - b) - 2$ qubits in the complement, leaving exactly one remaining spin degree of freedom. Consequently, the reduced density matrix of the state $|\psi_{+1,-1}\rangle$ of both the complement and A has at most rank 2. Furthermore, as explained in Fig. 3.4(d), these two states differ in their eigenvalues of $\prod_{\nu \sim \mathcal{L}_x, \nu \in \bar{A}} \sigma_\nu^z$. As a result, the distillable entanglement of an arbitrary superposition is zero.

The only symmetry in the transverse Ising model is the global spin-flip symmetry, which gives rise to *two* superselection sectors within which entangled superpositions of the scar states can exist, see Fig. 3.3 for odd L .

3.6 Generating Algebra

Here we discuss the commutant algebra that generates the scar subspace. As is shown in Appendix 3.C, the scar subspace is preserved not just by H , but also by the individual terms $X_p + X_{p+L}$, $Z_p Z_{p+1} + Z_{p+L} Z_{p+L+1}$, and $Z_p Z_{p+L}$ appearing in H . Following [170, 171] we denote the algebra generated by these terms as $\mathcal{A}$, and its commutant — that is, the algebra of all operators that commute with all operators in $\mathcal{A}$ in the endomorphism algebra of the full Hilbert space — as $\mathcal{C}$. The existence of QMBS is reflected in a non-trivial $\mathcal{C}$, in that one can find a basis in which all elements of $\mathcal{C}$ share the same block diagonal form, and the scar subspace corresponds to a particular block.

We compute the commutant $\mathcal{C}$ in the two leg transverse field Ising ladder by using a fermionic representation of the operator algebra introduced in Appendix 3.C. Here, we give a brief overview of the argumentation: this representation comes about from

¹By fixing the superselection sector, α and β , we create a state with a defect localized either inside or outside of A . A superposition with defect trapped *outside* of A would have all $X_p X_{p+L}$ and $Z_p Z_{p+L}$ identical everywhere *inside* of A .

applying the ordinary Jordan-Wigner (JW) transformation from left to right along the first row (sites $1, \dots, L$), and then along the second row (sites $L + 1, \dots, 2L$). The two-level fermionic system at each site $p = 1, \dots, 2L$ can be described using two Majorana operators $\gamma_p, \bar{\gamma}_p$, which can then be recombined into complex fermions $c_p, c_p^\dagger$ and $\bar{c}_p, \bar{c}_p^\dagger$, $p = 1, \dots, L$ by $c_p = \frac{1}{2}(\gamma_p - i\gamma_{p+L})$, $\bar{c}_p = \frac{1}{2}(\bar{\gamma}_p - i\bar{\gamma}_{p+L})$. The virtue of this representation is that two of the three types of terms generating $\mathcal{A}$ fermionize into manifestly particle number, $\hat{N} = \sum_{p=1}^L (c_p^\dagger c_p + \bar{c}_p^\dagger \bar{c}_p)$, conserving operators quadratic in the fermions: $X_p + X_{p+L} \rightarrow 2i(c_p \bar{c}_p^\dagger + c_p^\dagger \bar{c}_p)$ and $Z_p Z_{p+1} + Z_{p+L} Z_{p+L+1} \rightarrow 2i(\bar{c}_p c_{p+1}^\dagger + \bar{c}_p^\dagger c_{p+1})$ (for $p < L$). As shown in Appendix 3.C, the remaining terms, $Z_p Z_{p+L}$ and $Z_1 Z_L + Z_{L+1} Z_{2L}$, fermionize to operators that include a residual JW string, exhibiting a particle-hole symmetry that conjugates $\hat{N}$ to $2L - \hat{N}$. This implies that $(\hat{N} - L)^2$ commutes with these terms and, therefore, with all of $\mathcal{A}$, making it an element of $\mathcal{C}$. In fact, $(\hat{N} - L)^2$ generates the entirety of $\mathcal{C}$, as follows from Table II of [172]. In terms of the commutant $\mathcal{C}$, for the even L scar subspace the fermionic dual has overall even fermion parity, and the two scar states correspond to $(\hat{N} - L)^2 = L^2$, i.e. the all empty $\hat{N} = 0$ and the completely filled $\hat{N} = 2L$ states. For odd L the fermionic dual has odd fermion parity, and hence the least degenerate eigenspace of $(\hat{N} - L)^2$ corresponds to the eigenvalue $(L - 1)^2$, has dimension $4L$, and is spanned by the one-particle and one-hole states ².

3.7 Experimental Realization.

Finally, we propose an experimental implementation by time-evolving a scar stabilizer basis state in the LGT, versus preparing a non-scar state. Importantly, since the scar subspace is exactly decoupled from the rest of the Hilbert space, an initial scar

²The scars identified here differ from those in Ref. [173], where an exact low-entanglement scar state, $|\psi'_{E=0}\rangle$, was constructed as a tensor product over ‘diagonal’ singlets (Eq. 8 in Ref. [173]). This state, while an eigenstate of the two-leg transverse field Ising model, lies outside the scar subspaces computed in this work. In particular, the state $|\psi'_{E=0}\rangle$ is not robust against translation symmetry breaking in the x -direction, as opposed to the scar subspaces found in this work.

state is constrained to explore only a polynomially small subspace of $4L$ states, as opposed to a non-scar state, which may probe the exponentially large 2^{2L-1} dimensional Hilbert space. Focusing on odd L , the state preparation protocol for a simultaneous eigenstate of $Z_p Z_{p+L}$, $X_p X_{p+L}$, and the Gauss law operators is outlined in Fig. 3.5, with Fig. 3.5(a) showing a sample target state. This is achieved by configuring the horizontal and vertical links as depicted in Fig. 3.5(b), followed by the algorithm shown in Fig. 3.5(c) which shows that the scar basis states are 2-qubit product states. Non-scar initial states are randomly drawn (z -basis) product states that are consistent with Gauss law and have zero overlap with the scar subspace. We focus on two observables that distinctly differentiate scar dynamics from non-scar dynamics: the overlap between the initial and time-evolved state, known as the Loschmidt echo

$$L(t; |\psi_0\rangle) \equiv \langle \psi_0 | \exp\{-iHt\} | \psi_0 \rangle, \quad (3.12)$$

and the distillable von Neumann entropy defined in Eq. (3.8). While exact time evolution is performed here, in an experimental setting one would typically employ Trotterized time evolution—for an alternative scheme that avoids it see 3.H.

In Fig. 3.6(a), we plot the absolute value $|L(t)|$ as a function of time for one of the $4L$ scar basis states (solid red lines) for $L = 11$. For comparison, we also show the dynamics of a non-scar initial state (solid blue lines). The evolution of scar versus non-scar states differ dramatically. In the polynomially small scar subspace the Loschmidt echo decays to $|L(t)| \sim 1/\sqrt{4L}$. However, no full revivals are observed, as the energies of eigenstates in the scar subspace are not equidistant. In contrast, non-scar initial states exhibit quantum ergodic behavior, where the initial state explores the (exponentially large) full Hilbert space that is available to it based on energy conservation³. Eventually, this state thermalizes, with the Loschmidt echo decaying to a much lower value. For comparison, we show $|L(t)| \sim 1/\sqrt{2^{2L-1}}$, the saturation

³This space is, roughly, the space spanned by eigenstates in the window $E \pm \Delta E/2$, where $E = \langle H \rangle$ and $\Delta E^2 \equiv \langle H^2 \rangle - \langle H \rangle^2$, and $\langle \cdot \rangle$ the initial state expectation value.




Figure 3.5: *QMBS Preparation*. (a) Example of a scar basis state specified by the eigenvalues of $X_p X_{p+L}$ and $Z_p Z_{p+L}$. (b) Initial state for the preparation algorithm in the LGT formulation ($V_y = 1$). Horizontal links are initialized in σ^z (electric) eigenstates, while vertical links are initially set to $|0\rangle$. (c) The state preparation circuit operates on pairs of vertically separated qubits located at ν and $\nu + \hat{y}$. An example is shown where the circuit prepares a state with $\sigma_\nu^x \sigma_{\nu+\hat{y}}^x = -1$ and $\sigma_\nu^z \sigma_{\nu+\hat{y}}^z = -1$; combined with the horizontal links, this yields a Gauss law $+1$ eigenstate. Minor circuit adjustments enable preparation of any combination of $\sigma_\nu^x \sigma_{\nu+\hat{y}}^x = \pm 1$ and $\sigma_\nu^z \sigma_{\nu+\hat{y}}^z = \pm 1$ states, allowing access to any scar subspace basis state. The circuit maintains product-state structure along L ; each state remains a 2-qubit product state.




Figure 3.6: *Experimental Realization*. (a) The absolute value of the Loschmidt echo, $|L(t)|$, is shown as a function of time, comparing an initial state within the scar subspace (red lines) to a thermalizing state (blue lines) for system size $L = 11$. The time-averaged value at late times is indicated by the dark red dashed line for the scar state. The scar echo decays as $|L(t)| \sim 1/\sqrt{4L}$ (gray dotted line), where $4L$ is the size of the scar subspace. In contrast, the thermalizing state (blue lines) explores a much larger Hilbert space. For a random circuit that fully explores this space, $|L(t)|$ would decay to $|L(t)| \sim 1/\sqrt{2^{2L-1}}$ (gray dotted lines), however, due to energy conservation, the thermalizing state's late-time average (dark blue dashed line) remains slightly above this value. (b) The distillable entanglement S_{vN}^{dist} of a time-evolved scar state (red) remains zero throughout. By contrast, for a thermalizing state (light blue), S_{vN}^{dist} increases over time until reaching a constant, non-zero saturation value (blue dashed line).

value in the absence of energy conservation. Loschmidt echos can be measured in experiment using ancilla-based interferometry [174, 175, 176].

Fig. 3.6(b) shows the distillable entanglement, revealing an even more striking difference between scar and non-scar sectors. In the scar sector, the distillable entanglement remains exactly zero at all times, consistent with the argument presented earlier. In contrast, for states outside the scar subspace, the distillable entanglement increases over time, indicating that the system is thermalizing. Using random-measurement tools [177], for the model at hand, (distillable) entanglement can be measured directly [178], or see [161] for an experimental demonstration via entanglement Hamiltonian tomography [179, 180].

3.8 Discussion

In this chapter, we identified exact quantum many-body scars in a quasi-1d limit of a 2+1 dimensional $\mathbb{Z}_2$ LGT, for arbitrary coupling, that are based on zero-magic resource stabilizer states. An Ising-LGT duality and mapping onto a fermionic representation allowed us to analytically show their existence via a non-trivial commutant of the algebra generated by the Hamiltonian terms which we explicitly compute. Notably, in the gauge theory formulation, scar states—and arbitrary superpositions—always have *zero distillable entanglement*. The analytical construction presented in this manuscript does not extend beyond the $L \times 2$ geometry. The specific scar states we identified are absent in larger geometries such as $L \times 3$ and $L \times 4$. Although we have explored these cases numerically for a limited range of coupling values g , constrained by the exponentially increasing computational cost, we cannot rule out the existence of other, qualitatively different types of QMBS in these geometries.

Our finding reveals a connection between computational complexity [100, 181, 182, 183, 184, 185, 186] and quantum ergodicity, finding that states with zero distillable entanglement also possess low non-stabilizerness. Additionally, it indicates that identifying QMBS through entanglement may be unreliable: If we had computed the

(total) von Neumann entropy, instead of the *distillable* entanglement, we likely would not have initially detected these states. In contrast, while not explicitly computed, the QMBS subspace can be distinguished from generic states, i.e., those that obey the ETH and exhibit quantum ergodicity, through measures of magic resource such as stabilizer Rényi entropies (SRE) [98]. In particular, we highlight [101], which investigates SREs of scarred states in the PXP model. Specifically, typical states have magic resource that scales *linearly* with L [98, 99], whereas the magic resource of states in the scar subspace is constrained by the superpositions that can be formed among the $4L$ zero-magic resource states. Based on arguments presented in [98], we conjecture that the magic resource of states within the scar subspace is at most *logarithmic* in L ⁴.

In the experimental realization we propose, classical simulability or intractability depends solely on the choice of the initial state, while the quantum circuit for time evolution remains fixed. This distinction is useful for verification and benchmarking, and is pursued in a forthcoming manuscript and Chapter 4.

Data availability— The data that support the findings of this chapter are openly available [187].

3.A Lattice Gauge Theory and Ising Dual

In this subsection, we provide details and nomenclature of the $\mathbb{Z}_2$ LGT model that we work with, along with the mapping to a $(2 + 1)$ d Ising model. The model described by Eq. (3.1) is well studied, as is the duality, and details can be found for instance in Refs. [167, 168, 169].

Eq. (3.1) consists of a 'magnetic term' (or 'plaquette term'), $\sum_p \prod_{\nu \in p} \sigma_\nu^x$, and an

⁴Our argument is based solely on the observation in [98] that, for a spin-1/2 (qubits) system, the α -SRE is bounded by $\log(d)$ (specifically, it is bounded by $\log(d + 1) - \log(2)$ for $\alpha = 2$), where $d(=4L)$ is the dimension of the space. Although the scar subspace is not merely a system of qubits and the assumptions of [98] do not directly apply, we would find a different scaling with L surprising

'electric term,' $\sum_\nu \sigma_\nu^z$. This nomenclature is borrowed from lattice quantum electrodynamics (LQED), to which the model is related when spin- $\frac{1}{2}$ degrees of freedom are generalized to N -dimensional ones, yielding a $\mathbb{Z}_N$ LGT. In the limit $N \rightarrow \infty$, this recovers LQED without matter. The model describes a *gauge theory* because of the presence of local constraints, *Gauss laws*,

$$G_s \equiv \prod_{\nu \in s} \sigma_\nu^z. \quad (3.13)$$

These operators commute with the Hamiltonian, i.e., $[H, G_s] = 0$ for all s , thereby defining superselection sectors. The sector where all $G_s = 1$ is typically referred to as the 'physical' sector. Continuing the analogy with electrodynamics, this is the sector without background charges, where electric field lines originate from (positive) electric charges—specifically, $\mathbb{Z}_2$ charges in our case [169]. For periodic boundary conditions, the model described by Eq. (3.1) has two additional superselection operators, $[H, V_x] = [H, V_y] = 0$, where V_x and V_y are defined in the main text as 'ribbons' of electric field operators that wind periodically around the x - and y -dimensions of the lattice, respectively [169].

The LGT-Ising duality was first described in [167], and we refrain from discussing its details beyond those pertinent to our study. The precise mapping between the LGT and the Ising dual can be found in Table 3.1. In the Ising dual, the Gauss law constraints are eliminated, and V_x and V_y become parameters in the Hamiltonian (see Eq. (3.2)), which take values ± 1 depending on the superselection sector in the LGT. For periodic boundary conditions in the LGT, the product of all plaquette operators is unity by construction. On the dual side, this condition must be enforced by requiring $\prod_p X_p = 1$.

3.B Entanglement Structure and Symmetries

In this subsection, we provide additional discussion of the symmetries and entanglement structure of the LGT model and its Ising dual. The (2+1)d $\mathbb{Z}_2$ LGT with



	$\mathbb{Z}_2$ LGT	Ising Dual
Hilbert space		
Magnetic	 $\prod_{\nu \in p} \sigma_{\nu}^x$	 X_p
Electric	 σ_{ν}^z	 $Z_p Z_{p+\hat{a}}$
Gauss laws	 $\prod_{\nu \in s} \sigma_{\nu}^z$	Identity
Parity	Identity	$\prod_p X_p = 1$
Stabilizers	 $\prod_{\nu \in p} \sigma_{\nu}^x \sigma_{\nu+L}^x$	 $X_p X_{p+L}$
	 σ_{ν}^z	 $Z_p Z_{p+L}$

Table 3.1: *Mapping Between the $\mathbb{Z}_2$ LGT and its Ising dual.* The degrees of freedom are spin- $\frac{1}{2}$ variables, which reside on the links ν in the LGT and on the faces of the plaquettes p in the dual formulation. The correspondence between magnetic, electric, and Gauss law terms is outlined. In the dual picture, the Gauss laws—originally local constraints in the LGT—is eliminated. The parity constraint in the dual corresponds to the global identity that the product of all plaquettes equals one in the LGT. Additionally, the mapping of the $Z_p Z_{p+L}$ and $X_p X_{p+L}$ stabilizers is presented.

periodic boundary conditions exhibits several symmetries, including local Gauss laws, Eq. (3.13), at each site or ‘vertex’ s , as shown in Fig. 3.7(a). Since $\prod_s G_s = 1$, these operators define 2^{2L-1} distinct superselection sectors of the model. We focus on the sector where $G_s = 1$ for all s , commonly referred to as the ‘physical sector’. Additionally, there are two ribbon operators, depicted in Fig. 3.7(b),

$$V_{x,y} \equiv \prod_{\nu \in \mathcal{L}_{x/y}} \sigma_\nu^z, \quad (3.14)$$




Figure 3.7: *Symmetries of the LGT.* (a) Shown are $2L$ Gauss laws, Eq. (3.13), at sites s , which are local symmetry operators that commute with the Hamiltonian H , thereby defining superselection sectors. (b) For periodic boundary conditions, the ribbon operators V_x and V_y , Eq. (3.14), defined along $\mathcal{L}_x$ and $\mathcal{L}_y$, respectively, specify additional superselection sectors.

which wind non-trivially around the two periodic directions of the lattice and commute with the Hamiltonian, i.e., $[H, V_x] = [H, V_y] = 0$. These symmetries result in a total of 2^{2L+1} superselection sectors. If a state is an eigenstate of these operators—meaning it is not a superposition of different sectors—then any bipartition into a subsystem A and its complement $\bar{A}$,

$$\rho_A \equiv \text{Tr}_{\bar{A}}(\rho), \quad (3.15)$$

will inherit a specific symmetry structure, as discussed in the main text. The remnant of a Gauss operator, shown in Fig. 3.4(a) of the main text, that consists of three of

four legs of the Gauss law operator that are inside A , at a boundary site s' (i.e. s' is the site just beyond the plaquette a or just before plaquette b) is denoted by

$$\prod_{\nu \sim s'} \sigma_\nu^z = \sigma_{\nu_{1/2}^{a/b}}^z, \quad (3.16)$$

where $\nu \sim s'$ means that only operators within A are included in the product. In the equality, $\nu_{1/2}^{a/b}$ refers to one of horizontal links just outside of the subsystem, just as shown in Fig. 3.4(a) of the main text, the identity is just a reorganization of Gauss law, $G_s = 1$. The argument made in the main text, that any superposition of scar basis states in the LGT for odd L have zero distillable entanglement, is based on σ_ν^z , where ν' is either of the horizontal links $\nu_{1/2}^{a/b}$ in Fig. 3.4(a), and thus $\prod_{\nu \in s'} \sigma_\nu^z$ labeling superselection sectors.

To illustrate better that these (four) operators label symmetry sectors, note that they commute with ρ_A ,

$$[\prod_{\nu \sim s'} \sigma_\nu^z, \rho_A] = [\sigma_{\nu_{1/2}^{a/b}}^z, \rho_A] = 0, \quad (3.17)$$

and can thus be simultaneously diagonalized, defining symmetry blocks. This can be understood as follows: ρ_A can be expressed as a sum of Pauli strings within A ,

$$\rho_A = \frac{1}{2^{d_A}} \sum_{O \in P_A} \langle O \rangle O, \quad (3.18)$$

where P_A are all Pauli strings with non-zero expectation value within A , and d_A is the number of spins in A . Any of the four operators on the l.h.s of Eq. (3.16) may be part of this sum and may have a non-zero expectation value. However, no other operator in A fails to commute with it because it is identical, by means of Gauss law, to the r.h.s of Eq. (3.16), an operator clearly outside of A . Thus ρ_A commutes with it, marking superselection sectors.

The same applies to the ribbon operator shown in Fig. 3.4(b) of the main text, defined as $V_x^A = \prod_{\nu \sim \mathcal{L}_x^A}$, where $\mathcal{L}_x^A = \mathcal{L}_x \cap A$. Since $V_x = V_x^A V_x^{\bar{A}}$, with $V_x^{\bar{A}}$ being the analogous ribbon operator in the complement of A , the fixed value of V_x implies, by a

similar argument, that V_x^A defines a symmetry sector of ρ_A . For the other ribbon V_y , note that, since the lattice is bipartitioned along y , it remains a symmetry operator for both the global state and ρ_A ; however, as it equals the product of two Gauss laws along y , it does not define an independent symmetry block. Although there are six symmetry operators in total (four Gauss laws and two ribbons), only three of them are independent, leading to $2^3 = 8$ non-zero symmetry blocks.

The entanglement structure of the dual Ising model is fundamentally different from that of the LGT, even though both predict identical expectation values for any operator when mapped between the two formulations. The dual formulation possesses only one symmetry: parity,

$$P_X = \prod_{p=1}^{2L} X_p = 1 \quad (3.19)$$

which corresponds to the requirement that the product of all plaquettes equals one in the LGT, a condition necessary for the LGT-to-Ising mapping under periodic boundary conditions. This induces a two-block symmetry structure in the reduced density matrix describing any bipartition in the Ising dual. As with the LGT, the parity of A , $P_X^A = \prod_{p \in A} X_p$ is a symmetry of ρ_A , thereby defining the symmetry blocks: To maintain $P_X = 1$ parity globally, the subsystem can have either even or odd X -parity, or be in any superposition thereof.

Given the specific entanglement structure arising from the symmetries of a subsystem in either formulation, it is natural to decompose the von Neumann entanglement entropy into two components: the ‘symmetry’ component, $S^{\text{symm.}}$, and the ‘distillable’ component, $S^{\text{dist.}}$, as defined in Eq. (3.7) of the main text. Given the number of symmetry sectors, the symmetry entanglement can be at most $\log(2)$ in the Ising formulation, and $3\log(2)$ in the LGT. As discussed in the main text and further elaborated below, the scar states in the LGT exhibit zero distillable entanglement (for even and odd L) but possess non-zero symmetry entanglement. This characteristic is important: without separating the von Neumann entropy into these components,

distinguishing scar states from non-scar states and revealing their stabilizer origin would have been impossible.

3.C Derivation of the Scar Solutions

The $\mathbb{Z}_2$ LGT model with periodic boundary conditions, Eq. (3.1) of the main text, can be mapped onto a two-leg transverse field Ising model (TFI²), Eq. (3.2) of the main text, under the constraint $\prod_p X_p = 1$ which ensures that the product of all plaquettes is one, an identity on the LGT side, see Fig. 3.8(a) and (b). The TFI² is mapped onto a fermionic system via a Jordan-Wigner transformation that runs from the first spin in the top row to the last spin in the top row, then continuing at the first spin on the bottom row and ending at the last spin in the bottom row, see Fig. 3.8(b). As shown in Fig. 3.8(c), this is written in terms of Majorana operators, $\gamma_p, \bar{\gamma}_p$, $p = 1, \dots, 2L$ with $\gamma_p^2 = \bar{\gamma}_p^2 = 1$ and $\{\gamma_p, \gamma_q\} = \{\bar{\gamma}_p, \bar{\gamma}_q\} = 2\delta_{pq}$, $\{\gamma_p, \bar{\gamma}_q\} = 0$, as follows

$$X_p = -i\gamma_p\bar{\gamma}_p, \quad Z_p = \prod_{q=1}^{p-1} (i\gamma_q\bar{\gamma}_q)\gamma_p. \quad (3.20)$$

Next, we define two species fermionic operators by combining the Majorana modes *vertically along y*, i.e.,

$$\begin{aligned} c_p &= \frac{1}{2}(\gamma_p - i\gamma_{p+L}), & \bar{c}_p &= \frac{1}{2}(\bar{\gamma}_p - i\bar{\gamma}_{p+L}), \\ c_p^\dagger &= \frac{1}{2}(\gamma_p + i\gamma_{p+L}), & \bar{c}_p^\dagger &= \frac{1}{2}(\bar{\gamma}_p + i\bar{\gamma}_{p+L}). \end{aligned} \quad (3.21)$$

The fermion parity operators are given as

$$(-1)^{c_p^\dagger c_p} = i\gamma_p\gamma_{p+L}, \quad (-1)^{\bar{c}_p^\dagger \bar{c}_p} = i\bar{\gamma}_p\bar{\gamma}_{p+L}, \quad (3.22)$$

so that

$$X_p X_{p+L} \equiv -(-1)^{c_p^\dagger c_p + \bar{c}_p^\dagger \bar{c}_p}, \quad (3.23)$$

marks the combined-species fermion parity. We define the vacuum state by requiring $c_p|0\rangle = \bar{c}_p|0\rangle = 0$ for all p , with the mode ordering from left to right in ascending




Figure 3.8: *Analytic Derivation.* (a) The derivation of the scar eigenstates presented in this work begins with the degrees of freedom in the $\mathbb{Z}_2$ lattice gauge theory (LGT), represented by spins located on the links of a $L \times 2$ -site lattice with periodic boundary conditions. (b) This system is then mapped onto a dual Ising model, subject to the parity constraint $\prod_p X_p = 1$, where spin-1/2 degrees of freedom reside at the centers of the plaquettes (denoted by $p = 1, \dots, 2L$). A Jordan-Wigner transformation (indicated by the green arrow) initiates in the top row (covering plaquettes $1, \dots, L$) and extends to the bottom row (from sites $L + 1$ to $2L$), transforming the system into a two-component Majorana representation. (c) The Majorana modes are combined vertically to form a two-flavor fermion representation, where $p = 1, \dots, L$ denotes one-dimensional sites. This representation enables the identification of scar subspaces for both even and odd L via the commutant of the algebra generated by the terms in H .

order, i.e., $c_1^\dagger \bar{c}_1^\dagger \dots c_L^\dagger \bar{c}_L^\dagger |0\rangle$. Note that the constraint

$$1 = \prod_{p=1}^{2L} X_p = \prod_{p=1}^L X_p X_{p+L} = (-1)^L \prod_{p=1}^L (-1)^{c_p^\dagger c_p + \bar{c}_p^\dagger \bar{c}_p}, \quad (3.24)$$

either represents the fermion parity for even L or the *negative* fermion parity for odd L . Consequently, for even L , only an even number of fermion or hole excitations are allowed, whereas for odd L , only an odd number of fermion or hole excitations are possible. This property is crucial for the scar solutions.

In this representation, two of the three terms of the Ising Hamiltonian, Eq. (2) of the main text, conserve the total fermion number,

$$\hat{N} = \sum_{p=1}^L (c_p^\dagger c_p + \bar{c}_p^\dagger \bar{c}_p), \quad (3.25)$$

which ranges from $\hat{N} = 0$ to $\hat{N} = 2L$. Concretely, using Eqs. (3.20-3.21), we can rewrite the first term in Eq. (3.2) of the main text as

$$\begin{aligned} X_p + X_{p+L} &= i\bar{\gamma}_p \gamma_p + i\bar{\gamma}_{p+L} \gamma_{p+L} \\ &= 2i(\bar{c}_p c_p^\dagger + \bar{c}_p^\dagger c_p). \end{aligned} \quad (3.26)$$

A similar derivation for the third term, for $p < L$, yields

$$\begin{aligned} Z_p Z_{p+1} + Z_{p+L} Z_{p+L+1} &= i\bar{\gamma}_p \gamma_{p+1} + i\bar{\gamma}_{p+L} \gamma_{p+L+1} \\ &= 2i(\bar{c}_p c_{p+1}^\dagger + \bar{c}_p^\dagger c_{p+1}), \end{aligned} \quad (3.27)$$

demonstrating that both terms conserve particle number. For $p = L$, however, we find

$$Z_1 Z_L + Z_{L+1} Z_{2L} = \prod_{q=1}^L (i\gamma_q \bar{\gamma}_q) (i\gamma_1 \bar{\gamma}_L + i\gamma_{L+1} \bar{\gamma}_{2L}) \quad (3.28)$$

where $i\gamma_1 \bar{\gamma}_L + i\gamma_{L+1} \bar{\gamma}_{2L} = 2i(c_1^\dagger \bar{c}_L + c_1 \bar{c}_L^\dagger)$ conserves particle number, but the string $\prod_{q=1}^L (i\gamma_q \bar{\gamma}_q)$ does not. This string arises because $Z_1 Z_L + Z_{L+1} Z_{2L}$ is a non-local

operator connecting across the periodic boundary condition, reflecting the choice of Jordan-Wigner map between spins and fermions shown in Fig. 3.8(b). Similarly, the remaining term in Eq. (3.2) of the main text also involves also a Jordan-Wigner string:

$$Z_p Z_{p+L} = i \prod_{q=0}^{L-1} (i \gamma_{p+q} \bar{\gamma}_{p+q}) i \gamma_p \gamma_{p+L} \quad (3.29)$$

which does not conserve fermion number either.

Despite not conserving particle number, both Eq. (3.28) and Eq. (3.29) exhibit particle-hole symmetry, i.e., conjugating $\hat{N}$ to $2L - \hat{N}$ and vice versa. Consequently, all terms in the Hamiltonian conserve $(\hat{N} - L)^2$. As we will show below, the scar solutions are to be found, for even L , in the $(\hat{N} - L)^2 = L^2$ subspace, corresponding to the completely empty, $\hat{N} = 0$, or fully filled, $\hat{N} = 2L$, states. For odd L , the scars lie in the $(\hat{N} - L)^2 = (1 - L)^2 = (2L - 1 - L)^2$ subspace, i.e., the combined single-particle, $\hat{N} = 1$, and single-hole, $\hat{N} = 2L - 1$, states.

Equipped with this, we will now explicitly demonstrate that Eq. (2) of the main text posses exact scar eigenstates for even L which are the Fock vacuum and the fully occupied states, and furthermore that for odd L the scar subspace is spanned by $4L$ single-particle and single-hole excitations of the $\bar{c}^\dagger$ and $c^\dagger$ modes, H^{dual} has no matrix elements connecting these to any other state.

Scar Solutions for even L

Given the properties discussed in the previous section, we posit that the two scar solutions for even L are found in the smallest eigenspaces of $(\hat{N} - L)^2$: the vacuum state $|0\rangle$ and the fully occupied state

$$|\Lambda\rangle \equiv \prod_p c_p^\dagger \bar{c}_p^\dagger |0\rangle \quad (3.30)$$

where the product is in ascending order. These are eigenstates with eigenvalue zero of all terms in Eq. (3.2) of the main text, and by extension of Eq. (3.1) of the main

text, which we will show now. Both $|0\rangle$ and $|\Lambda\rangle$ are uniquely specified via the fermion parity, specifically for these states

$$1 = (-1)^{c_p^\dagger c_p + \bar{c}_p^\dagger \bar{c}_p} = -X_p X_{p+L}, \quad (3.31)$$

where we used Eq. (3.23), and

$$\begin{aligned} 1 &= (-1)^{\bar{c}_p^\dagger \bar{c}_p + c_{p+1}^\dagger c_{p+1}} = \bar{\gamma}_p \gamma_{p+1} \bar{\gamma}_{p+L} \gamma_{p+L+1} \\ &= -Z_p Z_{p+1} Z_{p+L} Z_{p+L+1}, \end{aligned} \quad (3.32)$$

where in the last equality we have used that $\gamma_p \equiv \prod_{k=1}^{p-1} (-X_k) Z_p$ and $\bar{\gamma}_p \equiv -i X_p \times \prod_{k=1}^{p-1} (-X_k) Z_p$. From Eq. (3.31) it follows that the space spanned by $|0\rangle$ and $|\Lambda\rangle$ must have $X_p X_{p+L} = -1$ and from Eq. (3.32) that $Z_p Z_{p+L} = -Z_{p+1} Z_{p+L+1}$ for all p . We can thus pick the two states given in Eq. (3.3) of the main text to span the even- L subspace.

We will show now that these are *exact* eigenstates with eigenvalue zero for any value of the coupling g . Specifically, $\rho_{1/2}$ are eigenstates of, i.e. commutes with, every individual term in Eq. (3.3) of the main text. We write $H^{\text{dual}} \equiv H_A + H_B + H_C$ where

$$H_A \equiv - \sum_{p=1}^L [X_p + X_{p+L}], \quad (3.33)$$

$$H_B \equiv - \tilde{g} \sum_{p=1}^L Z_p Z_{p+L}, \quad (3.34)$$

$$H_C \equiv - \sum_{p=1}^L g_p [Z_p Z_{p+1} + Z_{p+L} Z_{p+L+1}]. \quad (3.35)$$

First note that if $X_p X_{p+L} = -1$ (Eq. (3.31)), then $X_p = -X_{p+L}$, and therefore $H_A |\Psi_{1/2}\rangle = 0$ (or alternatively $[H_A, \rho_{1/2}] = 0$). Furthermore, from Eq. (3.32) it also follows that

$$Z_p Z_{p+1} = -Z_{p+L} Z_{p+L+1} \quad (3.36)$$

and thus $H_C|\Psi_{1/2}\rangle = 0$. Finally, note that

$$H_B|\psi_{1/2}\rangle = -\tilde{g} \sum_{p=1}^L \pm(-1)^p |\psi_{1/2}\rangle = 0, \quad (3.37)$$

because L is even. Therefore $|\psi_{1/2}\rangle$ are degenerate exact eigenstates of H^{dual} with eigenvalue zero (and thus of Eq. (3.1) of the main text on the other side of the duality). Because the states $|\psi_{1/2}\rangle$, as well as $|0\rangle$ and $|\Lambda\rangle$, are product states, their von Neumann entanglement entropy is zero for any bipartition. More specifically, separate symmetry and distillable entanglement components can be defined via Eq. (3.7) of the main text. Both the distillable and the symmetry component are zero, $S_{\text{dist.}}^{\text{dual}} = S_{\text{sym}}^{\text{dual}} = 0$, for the even- L scars in the Ising formulation. In the LGT formulation the distillable entanglement is also zero, $S_{\text{dist.}}^{\text{LGT}} = 0$, but the scars are not product states. A non-zero symmetry component stems from the Gauss laws at the boundary separating a subsystem from its complement, see the previous discussion in Appendix 3.B.

Scar Subspace for odd L

For odd L , we now discuss a subspace spanned by scar states. The states $|0\rangle$ and $|\Lambda\rangle$ are not allowed for odd L because of Eq. (3.23), instead only states with an odd fermion-parity are permitted. We posit that the scar subspace is spanned by the following $4L$ single-particle and single-hole states,

$$|\psi_p\rangle \equiv c_p^\dagger |0\rangle, \quad |\bar{\psi}_p\rangle \equiv \bar{c}_p^\dagger |0\rangle, \quad (3.38)$$

$$|\chi_p\rangle \equiv c_p |\Lambda\rangle, \quad |\bar{\chi}_p\rangle \equiv -\bar{c}_p |\Lambda\rangle, \quad (3.39)$$

$|\Lambda\rangle$ is defined in Eq. (3.30) and the fermion-parity is consistent with Eq. (3.23) for odd L . The terms in the Hamiltonian can be written in terms of Majorana operators and, subsequently, via fermion Fock operators. The plaquette term is

$$H_A = - \sum_{p=1}^L (i\bar{\gamma}_p \gamma_p + i\bar{\gamma}_{p+L} \gamma_{p+L}) = -2i \sum_{p=1}^L (\bar{c}_p^\dagger c_p - c_p^\dagger \bar{c}_p), \quad (3.40)$$

which conserves the total fermion number, Eq. (3.25). The electric terms in the Hamiltonian can be written as

$$H_B = \tilde{g} \sum_{p=1}^L \prod_{q=0}^{L-1} (i\gamma_{p+q}\bar{\gamma}_{p+q}) \gamma_p \gamma_{p+L} \quad (3.41)$$

and

$$\begin{aligned} H_C = & -g \sum_{p=1}^{L-1} (i\bar{\gamma}_p \gamma_{p+1} + i\bar{\gamma}_{p+L} \gamma_{p+L+1}) \\ & -gV_x \prod_{q=1}^L (i\gamma_q \bar{\gamma}_q) (i\gamma_1 \bar{\gamma}_L + i\gamma_{L+1} \bar{\gamma}_{2L}) \end{aligned} \quad (3.42)$$

whose properties are more difficult to discern. To understand the action of H_C (Eq. (3.42)) note that

$$i\bar{\gamma}_p \gamma_{p+1} + i\bar{\gamma}_p \gamma_{p+1} = 2i(\bar{c}_p^\dagger c_{p+1} - c_{p+1}^\dagger \bar{c}_p), \quad (3.43)$$

$$i\gamma_1 \bar{\gamma}_L + i\gamma_{L+1} \bar{\gamma}_{2L} = 2i(c_1^\dagger \bar{c}_L - \bar{c}_p^\dagger c_1), \quad (3.44)$$

are particle number conserving terms. Importantly, the product in the second line of Eq. (3.42) can be written as

$$\prod_{q=1}^L (i\gamma_q \bar{\gamma}_q) = \prod_{q=1}^L i(c_q^\dagger \bar{c}_q^\dagger + c_q^\dagger \bar{c}_q + c_q \bar{c}_q^\dagger + c_q \bar{c}_q) \quad (3.45)$$

and it conserves the space spanned by Eqs. (3.38-3.39). Specifically, one can show that

$$\begin{aligned} \prod_{q=1}^L (i\gamma_q \bar{\gamma}_q) |\psi_p\rangle &= -i^L |\chi_p\rangle, \\ \prod_{q=1}^L (i\gamma_q \bar{\gamma}_q) |\bar{\psi}_p\rangle &= i^L |\bar{\chi}_p\rangle, \\ \prod_{q=1}^L (i\gamma_q \bar{\gamma}_q) |\chi_p\rangle &= i^L |\psi_p\rangle, \\ \prod_{q=1}^L (i\gamma_q \bar{\gamma}_q) |\bar{\chi}_p\rangle &= -i^L |\bar{\psi}_p\rangle. \end{aligned} \quad (3.46)$$

Moreover, for Eq. (3.41), one finds that $\gamma_p \gamma_{p+L} = 2i c_p^\dagger c_p$ and that

$$\begin{aligned}
\prod_{q=0}^{L-1} (i\gamma_{p+q} \bar{\gamma}_{p+q}) |\psi_p\rangle &= i^L (-1)^p \theta(k, q) |\chi_p\rangle, \\
\prod_{q=0}^{L-1} (i\gamma_{p+q} \bar{\gamma}_{p+q}) |\bar{\psi}_k\rangle &= -i^L (-1)^p \theta(k, q) |\bar{\chi}_k\rangle, \\
\prod_{q=0}^{L-1} (i\gamma_{p+q} \bar{\gamma}_{p+q}) |\chi_k\rangle &= -i^L (-1)^p \theta(k, q) |\psi_k\rangle, \\
\prod_{q=0}^{L-1} (i\gamma_{p+q} \bar{\gamma}_{p+q}) |\bar{\chi}_k\rangle &= i^L (-1)^p \theta(k, q) |\bar{\psi}_k\rangle,
\end{aligned} \tag{3.47}$$

where $\theta(k, q) \equiv 1$ if $k \leq q$ and -1 otherwise. This establishes that H^{dual} (and thus H) has no matrix elements between the scar and non-scar subspaces.

While a rather lengthy derivation brought these results, a more elegant solution can be found noting that Eqs. (3.38-3.39) are uniquely specified by

$$(-1)^{c_k^\dagger c_k + \bar{c}_k^\dagger \bar{c}_k} = \begin{cases} 1 & \text{if } k = p \\ -1 & \text{otherwise} \end{cases}, \tag{3.48}$$

specifying the values of $X_p X_{p+L}$, and

$$(-1)^{\bar{c}_p^\dagger \bar{c}_p + c_{p+1}^\dagger c_{p+1}} = \begin{cases} -1 & \text{if } k = p \text{ or } k = p - 1 \\ 1 & \text{otherwise} \end{cases} \tag{3.49}$$

specifying the values of $Z_p Z_{p+L}$. The $4L$ states obeying Eqs. (3.48-3.49) are given in Eq. (3.5) of the main text. We will now work out the matrix elements of the Hamiltonian in these states. To do so, we first write them as

$$|\varphi_\alpha(k)\rangle \equiv \prod_{q \neq k} \left[\frac{1 - X_q X_{q+L}}{\sqrt{2}} \right] \frac{1 + X_k X_{k+L}}{\sqrt{2}} |r_\alpha(k)\rangle \tag{3.50}$$

where $|r_\alpha(k)\rangle$ are Z -eigenstates that realize the desired $Z_p Z_{p+L}$ stabilizer property (but are not $X_p X_{p+L}$ eigenstates). Choosing a convention, where the spin in the

bottom row is always up, the spin in the top row selects the $Z_p Z_{p+L}$ eigenvalue, for instance for $\alpha = 1$,

$$|r_1(k)\rangle \equiv \prod_{p < k} \left| \begin{array}{c} s^{(p,k)} \\ \uparrow \end{array} \right\rangle \left| \begin{array}{c} \overbrace{\downarrow}^k \\ \uparrow \end{array} \right\rangle \prod_{p > k} \left| \begin{array}{c} \bar{s}^{(p,k)} \\ \uparrow \end{array} \right\rangle, \quad (3.51)$$

where $s^{(p,k)} = \uparrow$ if $k - p$ is even and otherwise $\downarrow$, and $\bar{s}^{(p,k)} = \downarrow$ if $k - 1$ is even and otherwise $\uparrow$. The cases $\alpha = 2, 3, 4$ follow directly from Eq. (3.5) of the main text. The matrix elements of H^{dual} Eq. (3.2) of the main text with individual terms given in Eqs. (3.33-3.33), and identically for the LGT Hamiltonian Eq. (1) of the main text, are now easily derived. For the plaquette term, $H_A = \sum_p (X_p + X_{p+L})$, the scar-subspace matrix elements read

$$\begin{aligned} \langle \varphi_4(k) | H_A | \varphi_1(k) \rangle &= \langle \varphi_1(k) | H_A | \varphi_4(k) \rangle \\ &= \langle \varphi_3(k) | H_A | \varphi_2(k) \rangle = \langle \varphi_2(k) | H_A | \varphi_3(k) \rangle = 2, \end{aligned} \quad (3.52)$$

while the scar-subspace matrix elements of $H_B = \sum_p Z_p Z_{p+L}$ are

$$\begin{aligned} \langle \varphi_1(k) | H_B | \varphi_1(k) \rangle &= \langle \varphi_3(k) | H_B | \varphi_3(k) \rangle = -1, \\ \langle \varphi_2(k) | H_B | \varphi_2(k) \rangle &= \langle \varphi_4(k) | H_B | \varphi_4(k) \rangle = +1. \end{aligned} \quad (3.53)$$

Finally, the matrix elements of $H_C = \sum_p (Z_p Z_{p+1} + Z_{p+L} Z_{p+1+L})$ are

$$\begin{aligned} \langle \varphi_3(k-1) | H_C | \varphi_1(k) \rangle &= \langle \varphi_1(k+1) | H_C | \varphi_3(k) \rangle \\ &= \langle \varphi_4(k-1) | H_C | \varphi_2(k) \rangle = \langle \varphi_2(k+1) | H_C | \varphi_4(k) \rangle = 2. \end{aligned} \quad (3.54)$$

While states are represented differently in the LGT, see the main text, these matrix elements are identical between the formulations.

3.D Details of the Quantum Many-body Scar Solutions


In this section of the Appendix we provide additional numerical details for the QMBS solution discussed in the main text, both in the LGT and the Ising formulation. All

results are for $V_x = V_y = 1$, though while the $V_y = -1$ sector is trivial, the QMBS solutions exist in all sectors.

Fig. 3.9 shows the distillable entanglement, for even $L = 8$ and odd $L = 7$, across a range of couplings, $g \in \{0.1, 0.5, 1.5, 2.5\}$. The scar solutions persist for any value of the coupling and, moreover, for the LGT have exactly zero entanglement (which we verified to within machine precision). In contrast, in the Ising dual the distillable entanglement is non-zero and, depending on the coupling, the scar solutions can be indistinguishable from typical eigenstates. Notably, while the limites $g \rightarrow 0$ and $g \rightarrow \infty$ are well understood for the ground states of the model, mid-spectrum states behave highly non-trivial in these limits with the scar states persisting. Additionally, in Fig. 3.9b we observe an intricate band structure within the typical eigenstates. We speculate that these are due to additional protected subspaces potentially those that, in the fermionic formulation, are singled out by $(\hat{N} - L)^2$, i.e., the combined k -particle and -hole spaces where $k = 2, \dots$, which an extension of our investigation may uncover.


Fig. 3.10 shows the different components of the von Neumann entanglement entropy, the distillable and the symmetry part. Shown are the distillable entanglement $S_{\text{vN}}^{\text{dist}}$ in the top row, the symmetry entanglement $S_{\text{vN}}^{\text{symm}} \equiv -\sum_s p_s \log(p_s)$ in the middle row, and the total von Neumann entropy $S_{\text{vN}} = S_{\text{vN}}^{\text{dist}} + S_{\text{vN}}^{\text{symm}}$ in the bottom row. The QMBS states are highlighted in red, all results are for $L = 7$ (odd) and $L = 8$ (even) lattice sizes, with coupling $g = 0.9$. Although the total von Neumann entropy of scar states remains somewhat distinguishable at this coupling, it tends to blend scar states with non-scar states at other couplings, potentially obscuring the scar solutions. This underscores the value of studying entanglement structure—rather than only entanglement entropy — as a more revealing marker to differentiate scar from non-scar states. Our analytical analysis shows that, because of their stabilizer origin, it is magic resource that truly differentiates scar from non-scar states.

To emphasize this point more, Fig. 3.11 shows the (total) von Neumann entropy




(a) Results for $g = 0.1$, close to the integrable point $g = 0$.


Figure 3.9: *Distillable Entanglement for Varying Couplings*. The following plots show $S_{\text{vN}}^{\text{dist}}$ for a range of couplings. Panels in the top row display results for an even system size $L = 8$, while the bottom row shows results for an odd system size $L = 7$. The left panels correspond to results from the LGT formulation, and the right panels to those from the Ising model. Scar states are highlighted with red stars. In the Ising dual formulation, scar states exhibit non-zero distillable entanglement, blending with other typical states and becoming indistinguishable. However, in the LGT, scar states maintain exactly zero distillable entanglement, which differentiates them clearly. While the limits $g \rightarrow 0$ and $g \rightarrow \infty$ yield ground states that are straightforward to understand, the behavior of mid-spectrum states is complex, with scar eigenstates persisting in these limits.



(b) Results for $g = 0.5$. We point out a band structure visible within the typical states. This structure may be due to other (exact or approximate) protected subspaces in the model.



(c) Results for $g = 1.5$.



(d) Results for $g = 2.5$.



(a) Results for odd system size $L = 7$.

Figure 3.10: *Contributions to von Neumann Entanglement Entropy.* The top row shows the distillable entanglement S_{vN}^{dist} , the middle row the symmetry entanglement $S_{vN}^{\text{symm}} \equiv -\sum_s p_s \log(p_s)$, and the bottom row the total von Neumann entropy $S_{vN} = S_{vN}^{\text{dist}} + S_{vN}^{\text{symm}}$. The left column contains results from the LGT, while the right column shows results for the Ising dual. Quantum many-body scar (QMBS) states are highlighted in red, with the maximum possible symmetry entanglement indicated by horizontal gray dashed lines. For $g = 0.9$, these results show that QMBS states can be clearly distinguished by their distillable entanglement in the LGT, where they exhibit exactly zero values. By contrast, the total von Neumann entropy, although widely used, is finite and therefore less reliable for identifying QMBS states. Indeed, at other coupling values and in smaller systems, scar states may blend with typical states in terms of S_{vN} even within the LGT, rendering them indistinguishable.



(b) Results for even system size $L = 8$.




Figure 3.11: *Total von Neumann Entanglement Entropy*. Shown is the (total) Neumann Entropy S_{vN} for lattice size $L = 7$ with different couplings and with the scars, identified via their vanishing *distillable* entanglement, highlighted in red. (a) $g = 0.1$. (b) $g = 0.9$, (c) $g = 1.5$, (d) $g = 2.5$, Especially at high and low coupling, panels (a) and (d), the spectrum starts separating into energy bands and identifying the scars becomes impossible based on S_{vN} when not separated into distillable and symmetry components. For smaller sized systems, distinguishing scars from non-scars based on S_{vN} , instead of via the distillable entanglement, becomes impossible even at intermediate couplings.

S_{vN} for $L = 7$ and various coupling strengths, $g \in \{0.1, 0.9, 1.5, 2.5\}$. For both small and large couplings, scar and non-scar eigenstates become indistinguishable as the scar states merge into the continuum of non-scar states. This underscores the importance of identifying them through their distillable entanglement, which remains exactly zero.

Finally, we note a technical but important detail in our numerical analysis. For even L , the scar states are degenerate in energy. Since we are using exact diagonalization, our numerics typically return a linear superposition of the two scar states, which are analytically identified in the main text, and this superposition often has finite distillable entanglement. To resolve this, we employ a standard technique to break the

degeneracy by an infinitesimal amount, ensuring that the scar subspace remains intact while making the two scar solutions infinitesimally non-degenerate. Specifically, we perturb the coupling in front of $Z_p Z_{p+L}$ for a randomly chosen p by modifying the coupling as $g \rightarrow g + \epsilon$, where $\epsilon = 0.001$. In cases where analytic solutions are unavailable, an alternative method to handle this issue is described in [188].

3.E Other Lattice Geometries

As discussed in the conclusion of the main text, our analytical construction does not extend beyond the $L \times 2$ geometry. The specific scar states we identified are absent in larger geometries such as $L \times 3$ and $L \times 4$. Furthermore, we have explored these cases numerically for a limited range of coupling values g and lattice sizes because we are constrained by the exponentially increasing computational cost. One example is shown in Fig. 3.12. We have found no evidence of low entanglement states, but cannot strictly rule out the potential existence of other, qualitatively different types of QMBS in these geometries.

3.F Details of the Experimental Realization

In this section of the Appendix, we detail the state-preparation circuits used to generate scar basis states, and discuss additional details. Time evolution is implemented using standard Trotterized circuits, which are well-documented in the literature for this model and are not discussed here.

The experimental procedure discussed in the main text, for odd L , involves preparing a scar-subspace basis state versus a state outside this subspace, followed by time evolution. To prepare a scar basis state in the LGT a simple circuit, shown in Fig. 3.5 of the main text, is employed. The circuit prepares simultaneous $X_p X_{p+L} = \pm 1$ (represented by $\prod_{\nu \in p} \sigma_\nu^x \prod_{\nu \in p+L} \sigma_\nu^x$ in the LGT), $Z_p Z_{p+L} = \pm 1$ (corresponding to σ_ν^z in the LGT) and Gauss law eigenstates, with eigenvalues chosen according to any of the $4L$ basis states. The circuit does not introduce horizontal entanglement between




Figure 3.12: *Distillable Entanglement of Eigenstates for Lattice Geometries $L \times k$, where $k > 2$.* ((a) Left Panel): No low entanglement states are observed for 4×4 (here shown for $g = 0.9$), and for ((b) Right Panel): 5×3 . While our analytical construction clearly does not extend to these geometries, this numerical study does not strictly rule out the possibility of different types of QMBS in these geometries, for instance potentially appearing only near specific coupling values.

qubits, showing the rather trivial entanglement structure of the QMBS basis states; however, what is noteworthy is that any superposition of these states maintains zero distillable entanglement under time evolution.

An alternative approach, illustrated in Fig. 3.13(a-e), uses mid-circuit measurements [189, 190, 191] with constant-depth [192, 193, 194, 195]. This method is not strictly necessary: Although the model's ground state, the toric code, is long-range entangled for $g < g_c$, the scar subspace basis states that span the QMBS eigenstates are simple 2-qubit product states at any coupling (though the QMBS themselves are not). The circuit illustrated in panels (a-c) begins from a state that is an eigenstate of both $Z_p Z_{p+L}$ and $X_p X_{p+L}$, but not of the Gauss law operators. Conversely, the




Figure 3.13: *Alternative Circuit for QMBS Preparation Using Mid-Circuit Measurements.* (a) Initial state for the preparation algorithm in the LGT formulation ($V_y = 1$). Horizontal links are in σ^z (electric) eigenstates, and vertical links in σ^x (magnetic) eigenstates. Alternating $|0\rangle$ and $|1\rangle$ (or $|+\rangle$ and $|-\rangle$) yields eigenstates of any combination of stabilizers $Z_p Z_{p+L}$ and $X_p X_{p+L}$. (b) Mid-circuit measurement-based scheme for simultaneous Gauss law eigenstate preparation using L ancillas. (c) The circuit measures the combined σ^z parity of two neighboring vertical links via controlled- σ^z ($\Gamma(\sigma^z)$) operations and an ancilla a . The outcome dictates controlled-NOT operations on a vertical qubit, correcting Gauss laws in upper and lower rows, simultaneously. Control instructions classically depend on the desired target scar (via the horizontal links' state) and V_y sector. (d) A different starting point is shown for the scar state-preparation circuit, a simultaneous eigenstate of the stabilizer $Z_p Z_{p+L}$ and the Gauss law constraints. (e) This initial states is not eigenstates of an $X_p X_{p+L}$; to transform it into asimultaneous eigenstate, a circuit is used involving mid-circuit measurements and L ancillas. In a single shot, it measures $\sigma^x_\nu \sigma^x_{\nu'}$ values, where ν and ν' denote vertically stacked, vertically oriented links. From this the value of all $X_p X_{p+L}$ is determined. Based on the measurement outcome and the desired target state, a controlled σ^z operation is applied.

circuit shown in panels (d-e) starts from states that are simultaneous eigenstates of both $Z_p Z_{p+L}$ and the Gauss laws, though not of $X_p X_{p+L}$. As in the original method, no entanglement is introduced along the horizontal direction of the lattice, resulting in the same final state.

To create a state outside the scar subspace, a z -product state that is also a Gauss law eigenstate is initialized. This state will naturally have well-defined $Z_p Z_{p+L} = \pm 1$ eigenvalues. To ensure it has no non-zero overlap with the scar subspace, the $Z_p Z_{p+L} = \pm 1$ pattern is carefully chosen.

In this manuscript, we focus on the Loschmidt echo—defined as the overlap between the initial and time-evolved states—and the entanglement of the time-evolved states as key observables. The Loschmidt echo, $L(t)$, is a widely used metric to characterize the return probability of time-evolved states and can be measured via an interferometric scheme. This approach involves preparing an ancilla in a Hadamard superposition, using it to control the time-evolution operator, and performing measurements in the x - and y - bases to obtain the real and imaginary components of $L(t)$, respectively [174, 175, 176]. In our analysis, we consider the behavior of $|L(t)|$ at asymptotic times. On a digital quantum computer this would require Trotterization. Using large Trotter steps risks artificially enhancing the return probability by causing unphysical recurrence to the initial state. Thus very deep circuits are required, which is challenging for near-term devices. Alternatively, instead of its asymptote, one may consider the decay of $|L(t)|$ at early times which behaves as $1/\sqrt{gt}$ for an initial state within the scar subspace as shown in Fig. 3.14 for $L = 901$. Computing the same for sufficiently large L , is classically intractable for a non-scar initial state and is thus not shown here. However, the Loschmidt echo of a thermalizing system should decay exponentially [196].

The distillable entanglement of time-evolved states can be measured directly using symmetry-conscious k -designs that have been proposed in [178], including for the model that is considered here. This scheme enables the measurement of all required




Figure 3.14: *Loschmidt Echo in the Scar Subspace for Large System Sizes L* . The plot shows the absolute value of the Loschmidt echo, $|L(t)|$, evaluated for scar initial states with system sizes up to $L = 901$. The echo exhibits a power-law decay, $\sim (gt)^{-1/2}$, approaching an asymptotic value of $1/\sqrt{4L}$ at long times. The fit error is determined by varying the fit range and L .

quantities such as $p_s = \text{Tr}(\rho_A^{(s)})$, as well as k -fidelities, $\text{Tr}((\rho_A^{(s)})^k)$. From these, the distillable von Neumann entropy can be computed using $S_{\text{vN}}^{(s)} = -\lim_{k \rightarrow 1+} \frac{d}{dk} \text{Tr}((\rho_A^{(s)})^k)$ for sufficiently large k , typically $k \approx 3 - 4$ is sufficient [178].

3.G Local Observables

In previous QMBS studies, a more common observable are expectation values of local operators. Fig. 3.15 shows the behavior of local observables for both scar-subspace initial states and thermalizing states. Specifically, we consider the expectation value of electric operators, $\langle \sigma_\nu^z \rangle$, for horizontal (top panel) and vertical links (bottom panel), comparing scar dynamics (red and orange lines) versus non-scar dynamics (blue and light blue curves). Each panel shows two curves corresponding to horizontal and




Figure 3.15: *Local Observables*. The expectation values of electric operators, $\langle \sigma_\nu^z \rangle$, for links oriented horizontally (top panel) and vertically (bottom panel), contrasting scar initial states (red and orange lines) with thermalizing non-scar initial states (dark and light blue lines) for $L = 9$ and $g = 0.9$. For non-scar initial states, the expectation values thermalize to a constant, asymptotic value (indicated by dashed blue lines). In contrast, the expectation values within the scar subspace exhibit oscillatory behavior, albeit without perfect revivals, a consequence of the non-equidistant energy spacing of scar eigenstates. This limits the utility of local observables, as distinguishing these oscillations from finite-volume effects and experimental imperfections may prove challenging. We argue that observables such as the Loschmidt echo or the distillable entanglement, however, allow a stronger distinction between scar and non-scar dynamics.

vertical links at distinct lattice locations. Notably, while scar states exhibit somewhat more pronounced oscillations compared to thermalizing states—where local observables tend toward stationary values—no significant recurrences are observed in the local observables for scar-initial states. This lack of recurrence arises from the non-uniform energy spacing of the scar eigenstates. Because oscillations of similar magnitude may arise from finite-volume effects or experimental imperfections, for instance slight gate over/under rotations vary the value of g over space and time and additionally smoothen the curves, we believe that local observables offer less contrast between scar and non-scar dynamics than observables like the Loschmidt echo and distillable entanglement. Future work will include a stability analysis of the scar solutions against experimental imperfections and Trotter evolution.

3.H Random Circuit Evolution

The protocol discussed in the main text compares Hamiltonian evolution within the scar subspace to a thermalizing state outside of it. While valuable from a physics perspective for studying ergodic versus non-ergodic dynamics, it involves Trotterization, which requires very deep circuits to achieve thermalization. To demonstrate a quantum advantage, an alternative is based on shallower randomized circuits. Towards this end, we use the following circuit,

$$\mathcal{U} = \prod_{i=1}^s U(\boldsymbol{\alpha}_i, \boldsymbol{\beta}_i, \boldsymbol{\gamma}_i), \quad (3.55)$$

where s is the number of circuit layers, and

$$\begin{aligned} U(\boldsymbol{\alpha}_i, \boldsymbol{\beta}_i, \boldsymbol{\gamma}_i) &\equiv \exp\left\{-\sum_{p=1}^L \alpha_i(p) Z_p Z_{p+L}\right\} \\ &\times \exp\left\{-\sum_{p=1}^L \beta_i(p) [X_p + X_{p+L}]\right\} \\ &\times \exp\left\{-\sum_{p=1}^L \gamma_i(p) [Z_p Z_{p+1} + Z_{p+L} Z_{p+L+1}]\right\}. \end{aligned} \quad (3.56)$$




Figure 3.16: *Experimental Realization with a Random Circuit.* (a) Evolution of scar and non-scar states through random circuits, as defined in Eq. (3.55). The absolute value of the Loschmidt echo, $|L(s)|$, is shown as a function of the number of circuit steps s , contrasting an initial state within the scar subspace (red lines) with a state outside this subspace (blue lines) for a system size $L = 11$. Dashed lines represents late-time averages. For the scar state, the echo decays to $|L(s)| \sim 1/\sqrt{4L}$ (gray dotted line), where $4L$ is the dimension of the scar subspace. Conversely, the non-scar state decays to $|L(s)| \sim 1/\sqrt{2^{2L-1}}$ (gray dotted line), reflecting the lack of an energy conservation constraint. (b) The distillable entanglement, $S_{\text{vN}}^{\text{dist}}$, for scar versus non-scar initial states exhibits similar behavior to that shown in the main text for Hamiltonian evolution.

The angles $(\boldsymbol{\alpha}_i, \boldsymbol{\beta}_i, \boldsymbol{\gamma}_i) \equiv (\{\alpha_i(p)\}, \{\beta_i(p)\}, \{\gamma_i(p)\})$ are randomly chosen for every layer, thus the evolution is not constrained by energy conservation as in the main text. In the following we choose $\alpha_i(p) = \alpha_i$, $\beta_i(p) = \beta_i$, and $\gamma_i(p) = \gamma_i$, i.e., independent of p , and draw $(\alpha_i, \beta_i, \gamma_i)$ in every layer from a single-qubit circular unitary ensemble. Other randomizations are also possible.

As in the main text, initial states within the scar subspace remain confined to that subspace under randomized evolution, while states outside it, no longer constrained by energy conservation, explore the full Hilbert space. In Fig. 3.16, we present the same observables as in our main study: the Loschmidt echo in (a) and the distillable entanglement in (b), both as functions of circuit steps s . The stark contrast between scar and non-scar dynamics is clearly evident.

We note, however, that many randomized circuits may be classically simulable (see, e.g., [197]) and their computational hardness is an area of active investigation. We believe, therefore, that the thermalization example in the main text thus provides a stronger case for a demonstration of a quantum advantage.

Chapter 4

BENCHMARKING QUANTUM SIMULATION AT SCALE

A pre-print of work has been posted as “Benchmarking Quantum Simulation at Scale”, by Jeremy Hartse, Mohsin Raza, Shravan Shravan, Niklas Mueller, and Ivan H. Deutsch on the arXiv [198]

4.1 Introduction

The motivation for quantum computers is to address problems whose complexity lies beyond the reach of classical methods [28, 199, 29, 200]. Target areas include cryptography [201], optimization [202], sampling [203, 204, 205, 206], and quantum simulation ranging from chemistry and materials to high energy and nuclear physics [8, 10, 207, 150, 208, 209, 22, 12, 157].

Quantum simulation offers a compelling route to quantum advantage, yet establishing it rigorously remains challenging [210, 31, 211, 212, 104, 32, 213, 214, 215, 216]. Apart from experimental challenges, numerous classical algorithms already efficiently describe properties of many quantum systems [217, 1, 218, 219]. Moreover, noise can render quantum problems classically simulable [220, 221, 222, 223, 224, 225] and the issue is not only whether a computation is correct, but *how* correct it is. Finally, verifying a classically infeasible computation may be as difficult as simulating the computation itself. Verifying the correctness and advantage in analog quantum simulators is an even greater challenge [226, 227].

Early quantum supremacy demonstrations have been based on the computational complexity of sampling from certain hard distributions [30, 228, 229, 230, 231]. Benchmarking such unstructured problems, however, is challenging: the very property that

makes classical emulation difficult also renders verification exponentially costly. Problem agnostic benchmarks may be difficult to scale [107, 108, 232, 233, 234, 235, 236]. In contrast, we argue that quantum simulation applications could provide extra structure for verification.

We propose a method for verifying large-scale quantum simulations of nonequilibrium dynamics, aiming at the near-term regime where a broad quantum advantage may become achievable. We make use of models hosting quantum many-body scars (QMBS) [54, 55, 56, 57, 58], related to lattice gauge theory (LGT) models, which are key in many domains [237, 162, 163, 238, 164, 165].

The classical simulability of QMBS is not a new idea [54, 55]. Our approach leverages the structure of recently proposed *stabilizer scars* [149, 103, 239, 102], enabling not only efficient classical computation, but also guaranteeing efficient benchmarking that constrains the fidelity when simulating non-classically simulable, non-scar dynamics.

4.2 QMBS-Based Benchmarking

We consider Trotterized time evolution of a model with Hamiltonian H , represented by a quantum circuit U , and an input state $|\psi(0)\rangle$. The task is to verify the fidelity of the output of a circuit $\sigma = \mathcal{E}_U(|\psi(0)\rangle\langle\psi(0)|)$ computed on a device, relative to an (ideal) target state $\rho = U|\psi(0)\rangle\langle\psi(0)|U^\dagger$, where $\mathcal{E}_U$ represents the noisy version of the ideal unitary circuit U . The ideal circuit is neither a Clifford nor otherwise special circuit and is assumed, though not proven for the present application, to be classically hard for a generic input. Our protocol is compactly summarized in Fig. 4.1 with the following properties: (i) It is applied to models that host an exact, polynomially large QMBS subspace, spanned by stabilizer states, embedded in an exponentially large Hilbert space of typical states, with dynamics in the QMBS subspace classically tractable. (ii) The fidelity $F(\rho, \sigma)$ of any state simulated within the QMBS subspace relative to the target ρ can be efficiently measured experimentally using direct fidelity




Figure 4.1: *Benchmarking Quantum Simulations at Scale.* A protocol for benchmarking the fidelity of quantum simulations of non-equilibrium dynamics in regimes beyond classical simulability. We consider models which are non-integrable, for which typical states, $|\psi_{\mathcal{H}}\rangle$, thermalize. Such dynamics are not efficiently classically simulable, and benchmarking performance for such states is intractable. For models that host stabilizer scars, under consideration here, we can exploit the fact that for initial states in this subspace, $|\psi_S\rangle$ dynamics are classically simulable and we can perform direct fidelity estimation efficiently. Assuming local depolarizing noise, we demonstrate for a specific stabilizer scar model [149] that the fidelity of classically non-simulable states is approximately equal to that of the states restricted to the scar subspace, thus providing a method to verify performance for arbitrary inputs.

estimation [110, 111], with both the number of shots and observables scaling polynomially with system size. (iii) Typical (non-scar) inputs cannot be classically simulated or have their fidelity directly efficiently measured, however, with local depolarization noise, we show their fidelity is nonetheless as good as that of the scar states.

4.3 Efficient Direct Fidelity Estimation in a Stabilizer Scar Subspace

To prove (ii), we adapt the direct fidelity estimation (DFE) protocol of Ref. [110]. The fidelity between a pure target state ρ and its noisy experimental realization σ is

$$F(\rho, \sigma) = \text{Tr}(\rho\sigma) = \sum_{k=1}^{d^2} \chi_\rho(k) \chi_\sigma(k), \quad (4.1)$$

$\chi_\rho(k) = \text{Tr}[\rho W_k]/\sqrt{d}$, $d = 2^n$ for n qubits, and $W_k \in \mathcal{P}_n = \{\mathbb{I}, X, Y, Z\}^{\otimes n}$ is a n -qubit Pauli observable; X, Y, Z are Pauli matrices. To estimate Eq. (4.1), one draws, with probability $P_\rho(k) = \chi_\rho^2(k)$, samples $\{k_i\}_{i=1}^s$ at random. One measures the estimator $\chi_\sigma(k_i)/\chi_\rho(k_i)$ where, in our case, $\chi_\sigma(k_i)$ is obtained from a Pauli measurement from experiment and $\chi_\rho(k_i)$ classically computed. The average of this estimator is guaranteed to be ϵ -close to the true fidelity with probability $1 - \delta$ if $s = \lceil 1/(\epsilon^2 \delta) \rceil$ if the expectation value $\chi_\sigma(k_i)$ is known exactly [110]. In a quantum simulation, however, $\chi_\sigma(k_i)$ is estimated from m_i copies of the state, requiring

$$m_i = \left\lceil \frac{2 \log(2/\delta)}{|\text{Tr}[\rho W_{k_i}]|^2 s \epsilon^2} \right\rceil. \quad (4.2)$$

Eq. (4.2) ultimately controls the efficiency of the protocol. For a typical state, the situation is hopeless, since $|\text{Tr}[\rho W_{k_i}]|^2$ is exponentially small in the system size for all but very special observables.

Consider now (i), a QMBS subspace of dimension d_s , polynomial in qubit number n , spanned by an orthonormal set of stabilizer basis states, $\{\varphi_\alpha\}$. Let $\rho \in \mathcal{S}$, where $\mathcal{S} = \text{span}(\{\varphi_\alpha\})$, be pure. DFE is efficient for an arbitrary superposition $\rho \in \mathcal{S}$. To see this, we analyze the distribution of $|\text{tr}[\rho W_{k_i}]|^2$ for such states by considering the stabilizer Rényi entropy (SRE) [98]

$$M_\alpha \equiv \frac{1}{1-\alpha} \log(\xi_\alpha), \quad \xi_\alpha \equiv \frac{1}{d} \sum_{W_p \in \mathcal{P}_n} \text{Tr}[\rho W_p]^{2\alpha}. \quad (4.3)$$

Exploiting a result of Haug and Piroli [240] [their Eq. (11)] which bounds the stabilizer fidelity using SREs, and as we show in App 4.A, the stabilizer fidelity of the worst

$\rho \in \mathcal{S}$ is lower bounded by $1/d_s$, we find that the α -SRE

$$M_\alpha \leq \frac{2\alpha}{\alpha-1} \log(d_s), \quad \forall \alpha > 1. \quad (4.4)$$

For a pure state ρ , Eq. (4.3) specifies a probability distribution, $P_\rho(k_i)$. Eq. (4.4) implies a *polynomial* lower bound on the average of this distribution,

$$\xi_2 = \mathbb{E}_{k_i \sim P_\rho} [\text{Tr}[\rho W_{k_i}]^2] \geq \frac{1}{d_s^4}. \quad (4.5)$$

Since $\text{Tr}[\rho W_{k_i}]^2$ is a bounded random-variable, it is sub-gaussian by Hoeffding's lemma, therefore, it has an exponential tail-bound:

$$\Pr[|\text{Tr}[\rho W_{k_i}]^2 - \mathbb{E}_{k_i \sim P_\rho} [\text{Tr}[\rho W_{k_i}]^2]| \geq \gamma] \leq 2e^{-2\gamma^2}, \quad (4.6)$$

where $\gamma > 0$ [241], thus the distribution of Pauli expectations is concentrated around the inverse polynomial lower-bound; see Appendix (App) 4.A.

A concrete bound on the total number of copies, $N_\rho = \sum_{i=1}^s m_i$ can be obtained following Leone *et al.* [112]. Using an ϵ -truncation of the state [110], they derive an upper bound on N_ρ in terms of the 0-SRE. In $\mathcal{S}$, M_0 scales only as $\log(n^2)$ (more precisely, the set of nonzero expectation values is $d(n^2 - n + 2)$ for the model considered below) Consequently,

$$N_\rho^{[112]} \leq \frac{16}{\epsilon^4} \log(2/\delta) \mathcal{O}(d_s^2). \quad (4.7)$$

However, in practice no such ϵ -truncation needs to be introduced: An observable with exponentially small expectation value may occur but, because of Eqs. (4.5-4.6), a Markov Chain Monte Carlo (MCMC) algorithm for DFE (like the Metropolis-Hastings algorithm that we employ below) generating samples from P_ρ would accept those only with exponentially vanishing probability. Regardless of the argumentation, DFE is *efficient*: The fidelity between a state σ and an ideal pure state ρ can be efficiently estimated if $\rho \in \mathcal{S}$, the stabilizer scar subspace.

4.4 Fidelity Estimation for Classically Non-Simulable States

We now argue (iii), that under suitable assumptions on the noise, the fidelity estimate of a state in $\mathcal{S}$ also serves as a benchmark for the fidelity of non-classically computable non-scar states quantum computed with the same circuit. We assume, for simplicity, that the dominant error is approximated by local depolarizing noise with probability p , i.e., $\mathcal{E}(\rho) = \bigotimes_{i=1}^n \mathcal{E}^{(i)}(\rho)$ with $\mathcal{E}^{(i)}(\rho) = (1 - p)\rho + p\frac{\mathbb{I}}{2}$ and consider the *average* fidelity,

$$F(\mu) \equiv \mathbb{E}_{|\psi\rangle \sim \mu} [\langle \psi | \mathcal{E}(|\psi\rangle\langle\psi|) | \psi \rangle], \quad (4.8)$$

where the average is over input states drawn from a given ensemble μ , either the Haar measure over the full Hilbert space of n qubits, $|\psi\rangle \sim \mu_{\mathcal{H}}$, or the uniform measure over the stabilizer scar subspace, $|\psi\rangle \sim \mu_{\mathcal{S}}$, with dimension $d_s = \text{poly}(n)$. In the full Hilbert space,

$$F(\mu_{\mathcal{H}}) = \frac{1}{1 + 1/d} \left[\frac{1}{d} + \left(1 - \frac{3p}{4}\right)^n \right] \approx \left(1 - \frac{3p}{4}\right)^n \quad (4.9)$$

where, in the second relation, exponentially suppressed terms proportional to inverse powers of d have been neglected for $p < 1$ [$F(\mu_{\mathcal{H}}) = 1/d$ at $p = 1$]; see App 4.C for details. In contrast, for input states uniformly drawn from the scar subspace,

$$F(\mu_{\mathcal{S}}) = \frac{1}{d_s(d_s + 1)} \sum_{i=1}^{d_s^2} \left[\text{Tr}[K_i^{\mathcal{S}} K_i^{\mathcal{S}\dagger}] + |\text{Tr}[K_i^{\mathcal{S}}]|^2 \right], \quad (4.10)$$

where $K_i^{\mathcal{S}} \equiv \Pi_{\mathcal{S}} K_i \Pi_{\mathcal{S}}$ are the Kraus operators of the channel, proportional to Pauli operators W_i , i.e., $K_i \in \{(1 - \frac{3p}{4})^{\frac{1}{2}} \mathbb{I}, (\frac{p}{4})^{\frac{1}{2}} X, (\frac{p}{4})^{\frac{1}{2}} Y, (\frac{p}{4})^{\frac{1}{2}} Z\}^{\otimes n}$, projected onto $\mathcal{S}$ via $\Pi_{\mathcal{S}}$. We now derive a quantitative relation between the fidelity in the scar subspace, Eq. (4.10), and in the full Hilbert space, Eq. (4.9).

4.5 Prototype Model

With knowledge of error model and the structure of the scar subspace, Eq. (4.10) can be tightly lower and upper bounded. The following Ising model on $n = 2L$ qubits

posses a stabilizer scar subspace [149], see App 4.B, and satisfies both (i) and (ii),

$$H = - \sum_p X_p - g \sum_{p,a=1,2} Z_p Z_{p+\hat{a}}, \quad (4.11)$$

where $p = (p_1, p_2)$ defined on a lattice with size $L \times 2$ and periodic boundary conditions (PBC), and subject to a parity constraint, $\prod_p X_p = 1$ resulting in Hilbert space dimension $d = 2^{2L-1} \approx 2^{2L}$; X_p and Z_p are Pauli matrices. For odd L , the model hosts a QMBS subspace of dimension $d_s = 2n = 4L$, spanned by an orthonormal set of stabilizer basis states, $\varphi_{\alpha,k} = |\varphi_{\alpha,k}\rangle\langle\varphi_{\alpha,k}|$,

$$\varphi_{\alpha,k} \equiv \prod_{q=1}^L \frac{1 + s_{q,\alpha}^{(k)} Z_q Z_{q+L}}{2} \frac{1 + r_{q,\alpha}^{(k)} X_q X_{q+L}}{2} \quad (4.12)$$

where the notation is such that q labels sites starting in the top and continuing in the bottom row, i.e., $q \equiv (q_1, 1)$ and $q + L \equiv (q_1, 2)$; and $r_{q,a}^k = -1$ for $k \neq q$ and 1 otherwise, and

$$\begin{aligned} s_{q,1}^{(k)} &= s_{q,4}^{(k)} = -s_{q,2}^{(k)} = -s_{q,3}^{(k)} = \sigma_q^{(k)} (-1)^{k-q} \text{ for } q \neq k, \\ s_{k,2}^{(k)} &= s_{k,4}^{(k)} = -s_{k,1}^{(k)} = -s_{k,3}^{(k)} = 1 \text{ for } q = k. \end{aligned} \quad (4.13)$$

where $\sigma_q^{(k)} = +1$ for $q < k$ and -1 for $q > k$. Eq. (4.11) is dual to $\mathbb{Z}_2$ LGT; see App 4.B for details, or Ref. [149]. The evolution of any state in the subspace and any expectation value can be efficiently classically computed.

4.6 Local Depolarization Errors are Scar-Subspace Agnostic.

We consider the subspace spanned by Eq. (4.12) and compute Eq. (4.10) for local depolarization error. The first and second terms in Eq. (4.10) can be written, respec-




Figure 4.2: *Noise Acting on Scar vs. Non-Scar States.* Comparison between the fidelity under a local depolarizing channel, averaged over Haar-random input from the full Hilbert space, $F(\mu_{\mathcal{H}})$ (solid lines), and bounds of the respective fidelity for scar states $F(\mu_{\mathcal{S}})$ (shaded bands, Eq. (4.15) and App 4.C). The difference, shown in the lower panel, vanishes as $n \rightarrow \infty$ or for small p , and (trivially) for $p \rightarrow 1$. The tighter bound is shown, and the crossover between where Eq. (4.15) is tighter and where the App result applies is marked by square markers.

tively,

$$\sum_{i=1}^{d^2} \text{Tr}[K_i^{\mathcal{S}} K_i^{\mathcal{S}\dagger}] = \frac{1}{d} \sum_{i=1}^{d^2} (1-p)^{w[W_i]} \text{Tr}[\Pi_{\mathcal{S}} W_i]^2, \quad (4.14a)$$

$$\sum_{i=1}^{d^2} |\text{Tr}[\Pi_{\mathcal{S}} K_i]|^2 = \sum_{i=1}^{d^2} \left(1 - \frac{3p}{4}\right)^{2L-w[W_i]} \left(\frac{p}{4}\right)^{w[W_i]} \times |\text{Tr}[\Pi_{\mathcal{S}} W_i]|^2 \quad (4.14b)$$

where w is the Pauli weight of W_i . The only operators with non-zero trace in $\mathcal{S}$ are tensor products of the stabilizer generators, $X_p X_{p+L}$, $Y_p Y_{p+L}$, and $Z_p Z_{p+L}$, all others are traceless. From Eqs. (4.14a-4.14b), using arguments outlined in App 4.C, we bound

$$\begin{aligned} \left(1 - \frac{3p}{4}\right)^n &\leq F(\mu_{\mathcal{S}}) \\ &\leq \frac{1}{1+d_s} + \frac{d_s}{1+d_s} \left[\left(1 - \frac{3p}{4}\right)^2 + 3\left(\frac{p}{4}\right)^2 \right]^{\frac{n}{2}}. \end{aligned} \quad (4.15)$$

The lower and upper bounds bracket $F(\mu_{\mathcal{H}})$ and coincide with it, and with each other, asymptotically in n or for small p . However, away from this regime, the upper bound Eq. (4.15) is greater than or equal to $F(\mu_{\mathcal{H}})$ by sub-leading-in- n contributions. Thus the fidelity of non-simulable states may fall below that of the scar benchmark for larger p , in which case it may not provide reliable evidence of quantum advantage. Worse, the upper bound Eq. (4.15) is loose at $p = \mathcal{O}(1)$. To obtain tight bounds beyond this “weak-noise” regime (large n or small p), we exploit the structure of the scar subspace outlined in App 4.C to upper and lower bound $|\text{Tr}[\Pi_{\mathcal{S}} W]|$ for the selected operators W ; these have no compact form and are not written in the main text.

Fig. 4.2 illustrates the tightness of these combined bounds, where the top panel compares the fidelity of states in the full Hilbert space $F(\mu_{\mathcal{H}})$ (solid lines) with that of states in the scar subspace $F(\mu_{\mathcal{S}})$ (shaded bands). The difference shown in the bottom panel vanishes as $p \rightarrow 0$ or large n , and trivially as $p \rightarrow 1$. The weak-noise

regime $p \lesssim 1/n$ is contrasted with the constant-noise regime $p = \mathcal{O}(1)$, and the transition between the two bounds is marked by square markers in the bottom panel; the tighter bound is always shown. The averaged fidelity of non-verifiable typical states is as good as that of states in the scar subspace. Our bounds hold *on average*, i.e., input states to the error channel are drawn from the respective Haar measure rather than generated through circuit evolution. Below, we show numerically that the fidelities nearly coincide for circuit evolution from scar and non-scar initial states.

4.7 Numerical Simulations

We present numerical results using a Markov Chain Monte Carlo (MCMC) sampling algorithm for DFE, Eq. (4.1). The algorithm proposes only Pauli operators that have nonzero matrix elements within the target scar subspace, i.e., from the set $P_S \equiv \{W_k | \exists \rho \in \mathcal{S} \text{ s.t. } \text{Tr}[\rho W_k] \neq 0\}$; there are $|P_S| = (n^2 - n + 2)d$ such operators for the prototype model Eqs. (4.11-4.12); the precise algorithm is discussed in App 4.D. Since our focus is on demonstrating this efficiency rather than the MCMC sampler itself, we use a simple uniform proposal distribution over P_S .

Fig. 4.3(a) shows that the DFE estimate is ϵ -close to the true fidelity with probability $1 - \delta$ if $s = \lceil 1/(\epsilon^2 \delta) \rceil$ (dashed black line) if the Pauli expectation values are known exactly [110]. We estimate the fidelity between two pure target states, $|\rho\rangle$ and $|\sigma\rangle$ where $|\rho\rangle$ is chosen Haar-randomly from the scar subspace, and $|\sigma\rangle$ is obtained via time evolution within the same subspace such that $|\sigma\rangle = e^{-iHt}|\rho\rangle$, with $gt = 0.5$. For fixed states, we run 500 independent MCMC chains, each with a burn-in of 10^3 iterations and a thinning interval of 100, collecting 2×10^4 samples per chain. We plot the mean estimate of δ over all chains, with error bands indicating the 95% Wilks confidence interval [242]. As the figure shows curves for $L = 11, 31, 51, 71$ ($n = 22, 62, 102, 142$ qubits), the accuracy of the estimate is independent of n .

Fig. 4.3(b) shows the corresponding analysis for finite-shot estimation. The worst-case shot cost, Eq. (4.7), (per sample) (dashed line) is compared with the hypothetical




Figure 4.3: *Numerical MCMC Demonstration.* (a) Failure probability to reach precision $\epsilon = 5 \cdot 10^{-2}$ of MCMC-based DFE between a target $\rho \in \mathcal{S}$ and a reference state $\sigma \in \mathcal{S}$ as a function of the number of sampled Pauli operators with exact measurements. Results shown for $L = 11, 31, 51, 71$ ($n = 22, 62, 102, 142$ qubits) compared with the bound $s = \lceil 1/(\epsilon^2 \delta) \rceil$ [110] are independent of system size when expectation values are known exactly. (b) Finite-shot estimation. For $\rho \in \mathcal{S}$, we show the maximal shot number $\langle N_\rho \rangle$ (dashed black) versus the optimal shot cost (points; solid black fit), attainable if exact expectation values were known (see App 4.D) for precision $\epsilon = 5 \cdot 10^{-2}$ and failure probability $\delta = 1 \cdot 10^{-2}$. Shown are the average shot number over $s = 10^4$ Paulis, $\langle N_\rho \rangle = N_\rho/10^4$. samples. Both bounds scale as $\sim d_S^2$.

optimal shot cost (circles with solid black fit), corresponding to the number of shots required from binomial statistics if the exact expectation values $\text{Tr}[W_{k_i}\sigma]$ were known; see App 4.D for details. Both worst and optimal cases exhibit the same polynomial scaling, $d_S^2 \sim n^2$.

Next, we focus on simulating noisy quantum circuits for σ . We classically emulate, using QISKIT [243], the preparation of an initial basis state and its subsequent (noisy) time evolution, see Fig. 4.4 for circuit and error model. We apply single qubit depolarizing noise after every gate in the time evolution circuit but not for state preparation or measurement. The initial state preparation circuit consists of $L = n/2$ two-qubit Bell-pair sub-circuits applied to a product state $|s\rangle$, where s is an n -qubit bitstring. For special choices of s , the circuit prepares QMBS basis states, while for generic bitstrings it produces orthogonal non-scar states. The evolution circuit may either be Trotterized time evolution circuit or a random circuit preserving the subspace.

For simplicity, we employ a random time evolution circuit for $n = 10$ qubits of the subspace preserving form,

$$U(\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}) \equiv \prod_k^\ell e^{-i \sum_{i=1}^L \alpha_{k,i} [Z_i Z_{i+1} + Z_{i+L} Z_{i+L+1}]} \\ \times e^{-i \sum_{i=1}^L \beta_{k,i} Z_i Z_{i+L}} e^{-i \sum_{i=1}^L \gamma_{k,i} [X_i + X_{i+L}]} \quad (4.16)$$

where ℓ is the number of evolution layers and the angles $\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}$ are randomly (uniformly in i) chosen for every layer; single X (two-qubit ZZ) rotations are shown in blue (red) in Fig. 4.4¹. Fig. 4.5(a) compares the fidelity of a circuit evolved state within the scar subspace, $F(\rho_S, \mathcal{E}(\rho_S))$, with that of a typical state outside the subspace, $F(\rho_{\mathcal{H}}, \mathcal{E}(\rho_{\mathcal{H}}))$. Here, ρ_S is obtained by evolving the basis scar state $\varphi_{0,1}$ [Eq. (4.12)] under the ideal circuit, while $\mathcal{E}(\rho_S)$ is the corresponding output of the same circuit in the presence of noise. The states $\rho_{\mathcal{H}}$ and $\mathcal{E}(\rho_{\mathcal{H}})$ are defined analogously for an initial state outside the scar subspace using the same circuit, the only

¹A Trotter time evolution circuit is obtained by setting $\alpha_{k,i} = \beta_{k,i} = \gamma_{k,i} = \delta t$ where δt is the time step.




Figure 4.4: *Circuit Simulation*. Quantum circuit used to simulate Eq. (4.11), consisting of state preparation, (random or Trotterized) time evolution, and DFE. Local depolarizing noise is assumed, except during state preparation and measurement. The state preparation circuit consists of $L = n/2$ two-qubit Bell-pair sub-circuits and is applied to an input computational basis product state $|s\rangle$, where s is an n -qubit bit-string. For special s , the circuit prepares QMBS basis states, while for all others it produces orthogonal non-scar states.

difference being the input string s . The bottom panel shows the difference which vanishes with circuit layers ℓ . Although limited to modest system size, the fidelities are nearly indistinguishable at $n = 10$, and both are very close to average analytical results, indicating that noise does not distinguish between states in the scar subspace versus the full Hilbert space. Along with our analytical bounds, this indicates that our protocol accurately estimates the fidelity of classically intractable states, up to




Figure 4.5: *Circuit Simulation*. (a) Classical emulation of the circuit in Fig. 4.4 for $n = 10$ qubits comparing fidelities as a function of p for different circuit depths ℓ , comparing scar basis state (dashed lines) with non-scar inputs (dotted lines). The results are compared with the analytical expression for the *average* fidelity, Eq. (4.9), modified by an effective error rate $p_{\text{eff}} = 1 - (1 - p)^\ell$ accounting for the ℓ layers (solid black lines). The bottom panel shows the difference $\Delta F \equiv F_S - F_{\mathcal{H}}$, decreasing with ℓ . (b) Noisy circuit simulation with $n = 10$ qubits, $\ell = 15$ layers and $p = 0.001$, where Pauli expectation values are estimated from finite-shot sampling. Top: Fidelity versus the number of MCMC samples, s . Circles show shot-sampled results (10^4 shots per Pauli observable), while solid lines correspond to exact Pauli expectation values. Two representative MCMC chains are shown out of 100. Bottom: Final fidelity precision at the end of shot-sampled MCMC chains ($s = 10^4$).

controlled corrections that vanish asymptotically. Moreover, using concentration of the Haar measure, one can show that the fidelity equivalence holds not just on average,

but for most states (see App 4.C for details).

In Fig. 4.5(b) we emulate a noisy experiment by simulating the circuit in Fig. 4.4, starting from a stabilizer-scar basis state and applying local depolarizing noise with strength $p = 0.001$ to a system of $n = 10$ qubits evolved for $\ell = 15$ layers. We generate 100 independent MCMC chains and estimate the fidelity using either exact Pauli expectation values (solid lines) or finite-shot measurements with 10^4 shots per Pauli observable (circles; error bars denote the propagated statistical uncertainty). The top panel shows two representative MCMC trajectories, while the bottom panel summarizes the fidelity precision achieved after $s = 10^4$ MCMC samples. Consistent with Fig. 4.3, the observed precision and failure probability substantially outperform the bound Eq. (4.7).

4.8 Discussion

We presented a protocol for benchmarking quantum simulations at scale. The approach rests on three key ingredients: (i) classically simulability of scar states, (ii) efficient fidelity estimation via the stabilizer structure of the subspace, and (iii) that their fidelity tightly constrains that of classically non-simulable states evolved under the same circuit. The key to this protocol is that, although the unitary dynamics differ strongly between the scar and thermalizing subspaces, local decoherence acts similarly in both. This allows us to benchmark performance by only considering simulable dynamics, and estimating fidelity in the non-simulable regime.

There are several caveats. First, (iii) has only been proven as an average over input states to the error channel. Numerical evidence, however, strongly suggests that it also holds for individual states computed by quantum circuits. Second, our analysis assumes local depolarizing noise, and it remains to be seen how the results extend to other error models. However, while one can certainly engineer either a subspace or an error model that violates (iii), generic noise is unlikely so specific as to preserve the commutant of a nontrivial, digitally simulated Hamiltonian, let alone of arbitrary

systematically constructible models. Finally, it is unclear whether local stochastic noise is an effective description for logical qubits. Although logical error models can be derived under simplifying assumptions [244, 245], recent work suggests the possibility of non-Markovian logical noise [246, 247]; for approximate decoding [248, 249] logical errors are even less understood. Many practical improvements are also possible, for instance our MCMC algorithm is not efficient.

A next step is the implementation of the protocol on quantum hardware, where achieving fidelities beyond those attainable by classical algorithms would constitute a demonstration of quantum advantage. Our goal is benchmarking discovery-scale quantum simulations of nonequilibrium dynamics and thermalization in structured models such as LGTs [250, 158, 159], with applications to high-energy and nuclear physics, condensed matter, and materials science.

4.A Fidelity Estimation in Scar-subspace

In this Appendix, we show that direct fidelity estimation (DFE) is efficient for states in the scar-subspace. We do this by bounding the α -SRE of any state $\rho \in \mathcal{S}$ (target state).

4.A.1 Upper-Bound on α -SRE of ρ

We start by noting that $P_\rho(k_i) = \text{Tr}[\rho W_{k_i}]^2/d$ is a probability distribution for a pure state. Using this we define the stabilizer Renyi entropy of $\rho \in \mathcal{S}$ as [98]:

$$M_\alpha \equiv \frac{1}{1-\alpha} \log(\xi_\alpha), \quad \xi_\alpha \equiv \frac{1}{d} \sum_{W_k \in \mathcal{P}_n} \text{Tr}[\rho W_k]^{2\alpha}. \quad (4.17)$$

Note that M_0 is defined as the logarithm of the number of Paulis with a non-zero expectation value. Consider also the stabilizer fidelity of the state $|\psi\rangle$ defined as [251]:

$$\mathcal{F}(|\psi\rangle) \equiv \max_{|\phi\rangle \in \text{STAB}} |\langle \phi | \psi \rangle|^2, \quad (4.18)$$

where the maximization runs over the set of all stabilizer states (STAB). We now show that the stabilizer fidelity of the target state $\rho := |\psi_f\rangle\langle\psi_f| \in \mathcal{S}$ is lower bounded in the following way:

$$\mathcal{F}(|\psi_f\rangle) = \max_{|\phi\rangle \in \text{STAB}} |\langle\phi|\psi_f\rangle|^2 \geq \frac{1}{d_s}. \quad (4.19)$$

This is seen as following. First, note that given an orthonormal basis $S = \{x_1, x_2, \dots, x_{d_s}\}$ in $\mathbb{C}^{d_s}$, the maximum inner product between a given normalized vector $|\psi\rangle \in \mathbb{C}^{d_s}$ and any of the basis vectors is bounded by:

$$c_{\max}^2 = \max_i |\langle x_i | \psi \rangle|^2 \geq \frac{1}{d_s}, \quad (4.20)$$

where d_s is the dimension of $\text{span}(S)$. This bound holds even if we extend the set S with an arbitrary element $v \in \mathbb{C}^{d_s}$.

Take S to be the stabilizer basis of the scar subspace $\mathcal{S}$. The stabilizer fidelity of ρ is given by

$$\mathcal{F}(|\psi\rangle) \equiv \max_i |\langle\phi|\psi_f\rangle|^2, \quad (4.21)$$

where $|\phi\rangle_i \in \text{STAB} = \{|\phi\rangle_1, |\phi\rangle_2, |\phi\rangle_3, \dots, |\phi\rangle_m\}$, the set of all n -qubit stabilizer states. Since $|\psi_f\rangle \in \mathcal{S}$, the maximization over all stabilizer states in the RHS of Eq. (4.21) reduces to maximization over $\text{STAB} \cap \mathcal{S}$, which is an extension of the set S . Then, using Eq. (4.20), we get

$$\mathcal{F}(|\psi\rangle) \equiv \max_i |\langle\phi|\psi_f\rangle|^2 \geq \frac{1}{d_s}. \quad (4.22)$$

Now, we use the lower bound on the stabilizer fidelity to lower-bound the average value of $\text{Tr}[\rho W_{k_i}]^2$. This is done by using the connection between stabilizer fidelity with the α -SRE (for $\alpha > 1$) [240]:

$$M_\alpha(|\psi\rangle) \leq \frac{2\alpha}{\alpha - 1} \log[\mathcal{F}(|\psi\rangle)]^{-1}. \quad (4.23)$$

Combining Eq. (4.22) with Eq. (4.23) yields:

$$M_\alpha(|\psi\rangle) \leq \frac{2\alpha}{\alpha - 1} \log(d_s). \quad (4.24)$$

Thus, α -SRE is logarithmically upper-bounded for all states in the scar subspace for $\alpha > 1$. In Eq. (4.79) we show that the 0-SRE is upper-bounded logarithmically in system size.

4.A.2 Connection of α -SRE to DFE

Recall from Eq. (4.2) of the main text that the sampling cost of DFE to guarantee an ϵ - δ convergence is related to the minimum expectation value of the sampled Pauli W_{k_i} . The upper-bound in Eq. (4.24) implies that ξ_2 , which is the average value of random variable $\text{Tr}[\rho W_{k_i}]^2$ as seen in Eq. (4.17), is lower-bounded as:

$$\mathbb{E}_{k_i \sim P_\rho} [\text{Tr}[\rho W_{k_i}]^2] \geq \frac{1}{d_s^4}. \quad (4.25)$$

Consider the worst-case scenario where Eq. (4.25) is saturated. Since $\text{Tr}[\rho W_{k_i}]^2$ is a bounded random-variable, it is sub-gaussian by Hoeffding's lemma:

$$\Pr \left[\left| \text{Tr}[\rho W_{k_i}]^2 - \mathbb{E}_{k_i \sim P_\rho} [\text{Tr}[\rho W_{k_i}]^2] \right| \geq \gamma \right] \leq 2 \exp(-2\gamma^2), \quad (4.26)$$

where $\gamma > 0$ is a constant. Thus, it is exponentially unlikely for a Pauli expectation values to be significantly smaller than the average in Eq. (4.25). Combining Eq. (4.25), Eq. (4.26) and Eq. (4.2) immediately implies that DFE is efficient.

More rigorously, all Pauli expectation values that are smaller than Eq. (4.25) can be eliminated using the ' ϵ -truncation trick', wherein, one pays a small bias as a cost for the truncation (as seen in [110, 112]). This is done by constructing a truncated state ρ_2 , where any Pauli expectation with $|\text{Tr}\{\rho W_k\}| < \epsilon/d_s^2$ are ignored. Doing DFE with ρ_2 gives the correct fidelity up to an asymptotically vanishing bias: $|\text{Tr}\{\rho_2 \sigma\} - \text{Tr}\{\rho \sigma\}| \leq \epsilon/O(d_s)$. Since ρ_2 has bounded Pauli expectation values by definition, $|\text{Tr}\{\rho_2 W_k\}| \geq \epsilon/d_s^2$, the shot sampling cost of DFE is

$$\mathbb{E}_{k_i \sim P_{\rho_2}} [m_i] = \mathbb{E}_{k_i \sim P_{\rho_2}} \left[\frac{2 \log(2/\delta)}{\text{Tr}[\rho_2 W_{k_i}]^2 l \epsilon^2} \right] \leq \left\lceil \frac{2 d_s^4 \log(2/\delta)}{\epsilon^4} \right\rceil. \quad (4.27)$$

A tighter and more general version of this argument is presented in Ref. [112]. Considering total number of copies of the state $N_\rho = \mathbb{E}[m_i] = \sum_{i=1}^l m_i$, Leone et al. [112] show that

$$\frac{2}{\epsilon^2} \log(2/\delta) \exp[M_2] \leq N_\rho^{[112]} \leq \frac{64}{\epsilon^4} \log(2/\delta) \exp[M_0], \quad (4.28)$$

where M_2 and M_0 are the 2- and 0-SREs of ρ , respectively. Inserting in the upper-bound on M_0 from Eq. (4.79), directly leads to Eq. (4.7) of the main text.

In the main text we show that upper-bound derived from Eq. (4.28) scales as $\mathcal{O}(d_s^2)$ as compared to $\mathcal{O}(d_s^4)$ in Eq. (4.27). This may be a consequence of the looseness of Eq. (4.23). Numerical evidence presented in Ref. [240] suggests that the prefactor appearing in Eq. (4.23) may be smaller by a factor of half for $1 \leq \alpha \leq 2$, implying $\xi_2 \gtrsim d_s^{-2}$, in which case Eq. (4.28) and Eq. (4.27) scale identically.

4.B QMBS Model Details

In this Appendix, we provide additional details about the model, introduced in Ref. [149] and in Eq. (4.11) of the main text, and its QMBS subspace. The model is not unique in its hosting of stabilizer scars, and other models have been found or can be systematically constructed [103, 239, 102].

Details were previously presented in Ref. [149], and we provide only a high level summary. The model, Eq. (4.11), is defined on a rectangular $L \times 2$ lattice with periodic boundary conditions and parity constraint $\prod_{p=1}^{2L} X_p = 1$, in which case it is dual to $\mathbb{Z}_2$ LGT on $L \times 2$ plaquettes with Hamiltonian [167, 168, 169],

$$H_{\mathbb{Z}_2} = - \sum_p \prod_{l \in p} \sigma_l^x - g \sum_l \sigma_l^z, \quad (4.29)$$

where l labels a link of the dual lattice, and p labels a plaquette. The model is subject to $2L - 1$ independent Gauss law constraints, $G_s \equiv \prod_{\ell \in s} \sigma_\ell^z$, with $G_s |\psi\rangle = |\psi\rangle$ for “allowed/physical” states $|\psi\rangle$, where s labels a site in the dual lattice, thus matching the number of d.o.f. in Eq. (4.11) of the main text.




Figure 4.6: *Stabilizer scar subspace of the prototype model.* (a) The model, Eq. (4.11) of the main text, is defined on an $L \times 2$ lattice with periodic boundary conditions. The dual lattice of the LGT, Eq. (4.29), is shown in gray, in which the degrees of freedom reside on the links. (b) For odd L , the QMBS subspace is spanned by the orthonormal stabilizer basis states shown here, each characterized by a distinct eigenvalue pattern of the stabilizer generators $X_i X_{i+L}$ (blue) and $Z_i Z_{i+L}$ (red), where darker (lighter) shading indicates a +1 (−1) eigenvalue. (c) Initial states are prepared using simple Bell circuits specified by an input bit string s , which can generate any QMBS basis state or a state outside the scar subspace with the same circuit complexity. Subfigures (a) and (b) adapted from Ref. [149].

The stabilizer QMBS structure of the model originates from an underlying commutant algebra generated by fermionic operators. As shown in Ref. [149], the spin model, Eq. (4.11) of the main text, admits a Majorana representation via the Jordan-Wigner (JW) transformation. In this formulation, the Majoranas can be grouped

into different fermions, and the QMBS subspace corresponds to the conserved single-particle and single-hole subspace when L is odd. In this case, the operator $(\hat{N} - L)^2$, where $\hat{N}$ is the total fermion number operator, generates the full commutant of the scar subspace. In the spin model, the stabilizer generators $X_i X_{i+L}$ and $Z_i Z_{i+L}$ (for $i = 1, \dots, L$), see Fig. 4.6(a), correspond to relative parity operators between adjacent fermionic modes. The orthonormal stabilizer states spanning the QMBS subspace are depicted in Fig. 4.6(a) and (b), adapted from Ref. [149]. The matrix elements of the Hamiltonian within the scar subspace, which grows only linearly in the qubit number $n = 2L$, are derived in the Supplemental Material of that reference, enabling efficient simulation of scar dynamics for systems of arbitrary size.

State preparation of stabilizer-scar basis states is particularly straightforward. Fig. 4.6(c) illustrates the circuit used to prepare any such state. The circuit consists of a simple Bell-state preparation circuit acting on an input bitstring s , which is generated by applying NOT (X) gates to the corresponding qubits. Certain input bitstrings produce basis states of the form shown in Fig. 4.6(b), corresponding to stabilizer-scar states, while all other bitstrings generate non-scar states. Importantly, the preparation of scar and non-scar states requires exactly the same circuit and therefore has identical circuit complexity.

When simulating noisy circuits, we employ the QISKIT simulator [243] with a local depolarizing noise model of single-site error probability p , as illustrated in Fig. 4.4(a). The circuit is left uncompiled (relative to QISKIT's native gateset), so that the noise channel is inserted immediately after each gate layer as it appears in the figure. The elementary gate primitives are single-qubit rotations $\exp\{i\alpha_i X_i\}$ and two-qubit entangling gates $\exp\{i\alpha_{ij} Z_i Z_j\}$. The local depolarizing channel is applied after each such gate; no noise acts on the initial-state preparation circuit shown in Fig. 4.6(c), nor on the final measurement layer.

4.C Average fidelity of the noisy channel

In this Appendix, we analyze the fidelity of states within the scar subspace, and compare it with those in the full Hilbert space, under local depolarizing noise. For simplicity, we focus on the average channel fidelity with respect to input states drawn from a given measure μ , as defined in Eq. (4.8) of the main text. In particular, $\mu_{\mathcal{H}}$ denotes the Haar measure on the full Hilbert space, while $\mu_{\mathcal{S}}$ denotes the Haar measure restricted to the (polynomially sized) stabilizer scar subspace.

For the Haar measure $\mu_{\mathcal{H}}$, this can be calculated exactly following Ref. [252]:

$$F(\mu_{\mathcal{H}}) = \frac{1}{d(d+1)} \left(d + \sum_{i=1}^{d^2} |\text{Tr}(K_i)|^2 \right), \quad (4.30)$$

where d is dimension of the Hilbert space, and K_i are the Kraus operators for the local depolarizing channel with probability p , i.e., $K_i \in \{(1 - \frac{3p}{4})^{\frac{1}{2}} \mathbb{I}, (\frac{p}{4})^{\frac{1}{2}} X, (\frac{p}{4})^{\frac{1}{2}} Y, (\frac{p}{4})^{\frac{1}{2}} Z\}^{\otimes n}$, where X, Y, Z are Pauli matrices. All K_i are traceless, except, $K_1 = (1 - \frac{3p}{4})^{n/2} \mathbb{I}$, therefore

$$\sum_{i=1}^{d^2} |\text{Tr}(K_i)|^2 = d^2 \left(1 - \frac{3p}{4} \right)^n, \quad (4.31)$$

and one obtains Eq. (4.9) of the main text. Similarly, we can consider the average channel fidelity over states in the scar subspace, under the same error channel,

$$\begin{aligned} F(\mu_{\mathcal{S}}) &= \sum_{i=1}^{d^2} \mathbb{E}_{|\psi\rangle \sim \mu_{\mathcal{S}}} \left[\text{Tr} \left(|\psi\rangle\langle\psi| K_i |\psi\rangle\langle\psi| K_i^\dagger \right) \right] \\ &= \sum_{i=1}^{d^2} \mathbb{E}_{|\psi\rangle \sim \mu_{\mathcal{S}}} \left[\text{Tr} \left(|\psi\rangle\langle\psi|^{\otimes 2} (K_i \otimes K_i^\dagger) \mathbb{F} \right) \right], \end{aligned} \quad (4.32)$$

where $\mathbb{F}$ is the swap operator. Since $|\psi\rangle\langle\psi| = \Pi_{\mathcal{S}} |\psi\rangle\langle\psi| \Pi_{\mathcal{S}}$ for states in $\mathcal{S}$, projectors

$\Pi_{\mathcal{S}}$ on the subspace can be introduced,

$$\begin{aligned}
& \sum_{i=1}^{d^2} \mathbb{E}_{|\psi\rangle \sim \mu_{\mathcal{S}}} \left[\text{Tr} \left(|\psi\rangle \langle \psi|^{\otimes 2} (\Pi_{\mathcal{S}} K_i \Pi_{\mathcal{S}} \otimes \Pi_{\mathcal{S}} K_i^{\dagger} \Pi_{\mathcal{S}}) \mathbb{F} \right) \right] \\
&= \sum_{i=1}^{d^2} \text{Tr} \left(\frac{\mathbb{I}_{\mathcal{S}} + \mathbb{F}_{\mathcal{S}}}{d_s(d_s + 1)} (\Pi_{\mathcal{S}} K_i \Pi_{\mathcal{S}} \otimes \Pi_{\mathcal{S}} K_i^{\dagger} \Pi_{\mathcal{S}}) \mathbb{F} \right) \\
&= \sum_{i=1}^{d^2} \frac{1}{d_s(d_s + 1)} \left\{ \text{Tr}(K_i^{\mathcal{S}} \otimes K_i^{\mathcal{S}\dagger} \mathbb{F} + K_i^{\mathcal{S}} \otimes K_i^{\mathcal{S}\dagger} \mathbb{F} \mathbb{F}_{\mathcal{S}}) \right\} \\
&= \sum_{i=1}^{d^2} \frac{1}{d_s(d_s + 1)} \left\{ \text{Tr}(K_i^{\mathcal{S}} K_i^{\mathcal{S}\dagger}) + |\text{Tr}(K_i^{\mathcal{S}})|^2 \right\}, \tag{4.33}
\end{aligned}$$

yielding Eq. (4.10) of the main text. We have abbreviated $K_i^{\mathcal{S}} \equiv \Pi_{\mathcal{S}} K_i \Pi_{\mathcal{S}}$ and $\mathbb{F}_{\mathcal{S}}$ and $\mathbb{I}_{\mathcal{S}}$ are the swap and identity operators in $\mathcal{S}$, respectively; we used that $\mathbb{F}$ commutes with operators of the form $\hat{O} \otimes \hat{O}$, and that $\mathbb{F} \mathbb{F}_{\mathcal{S}}$ acts like identity on states within the scar subspace, i.e., $(\Pi_{\mathcal{S}} \otimes \Pi_{\mathcal{S}}) \mathbb{F} \mathbb{F}_{\mathcal{S}} (\Pi_{\mathcal{S}} \otimes \Pi_{\mathcal{S}}) = \mathbb{I}_{\mathcal{S}}$.

We will now proceed to bound both terms. We start by noting that when $p = 0$, $F(\mu_{\mathcal{H}}) = F(\mu_{\mathcal{S}}) = 1$, as in the channel $\mathcal{E}$ becomes the identity operation. Similarly, when $p = 1$, the channel maps all states to the maximally mixed state, $\mathcal{E}(\rho) = \mathbb{I}/d$. Since $(1/d) \langle \psi | \mathbb{I} | \psi \rangle = 1/d$ for all pure states $|\psi\rangle$, it follows that $F(\mu_{\mathcal{H}}) = F(\mu_{\mathcal{S}}) = 1/d$. Hence, we want bounds on $F(\mu_{\mathcal{S}})$ that converge to $F(\mu_{\mathcal{H}})$ at $p = 0$ and $p = 1$. To achieve this, we construct two bounds. One is tight at $p = 0$ and the other is tight at $p = 1$. The final bound will be the intersection of these two bounds.

4.C.1 Bound tight near $p = 0$

Consider the first term in Eq. (4.33),

$$\sum_{i=1}^{d^2} \text{Tr}(K_i^{\mathcal{S}} K_i^{\mathcal{S}\dagger}) = \text{Tr} \left(\Pi_{\mathcal{S}} \sum_{i=1}^{d^2} K_i \Pi_{\mathcal{S}} K_i^{\dagger} \right) = \text{Tr}(\Pi_{\mathcal{S}} \mathcal{E}(\Pi_{\mathcal{S}})) \tag{4.34}$$

where $\mathcal{E}(\cdot)$ is the error channel. Using the Pauli decomposition of $\mathcal{E}(\Pi_{\mathcal{S}})$, and assuming local depolarizing noise, we find

$$\sum_{i=1}^{d^2} \text{Tr}(K_i^{\mathcal{S}} K_i^{\mathcal{S}\dagger}) = \frac{1}{d} \sum_{j=1}^{d^2} (1-p)^{w[W_j]} \text{Tr}(\Pi_{\mathcal{S}} W_j)^2 \quad (4.35)$$

where $w[W_j]$ is the weight of the Pauli string W_j . Without considering the structure of $\text{Tr}(\Pi_{\mathcal{S}} W_j)$, an upper bound for this term, loose but accurate at small p , is

$$\sum_{i=1}^{d^2} \text{Tr}(K_i^{\mathcal{S}} K_i^{\mathcal{S}\dagger}) \leq \frac{1}{d} \sum_{j=1}^{d^2} \text{Tr}(\Pi_{\mathcal{S}} W_j)^2 = \text{Tr}(\Pi_{\mathcal{S}}^2) = d_s. \quad (4.36)$$

The lower bound can be found by keeping only the K_1 Krauss operator in Eq. (4.34), i.e.,

$$\begin{aligned} \sum_{i=1}^{d^2} \text{Tr}(K_i^{\mathcal{S}} K_i^{\mathcal{S}\dagger}) &\geq \text{Tr}(K_1^{\mathcal{S}} K_1^{\mathcal{S}\dagger}) \\ &= \text{Tr}(\Pi_{\mathcal{S}}) \left(1 - \frac{3p}{4}\right)^n = d_s \left(1 - \frac{3p}{4}\right)^n. \end{aligned} \quad (4.37)$$

Notice that the above two bounds are exact at $p = 0$, and hence tight. The second term in Eq. (4.10) of the main text can be written

$$\begin{aligned} \sum_{i=1}^{d^2} |\text{Tr}(K_i^{\mathcal{S}})|^2 &= \sum_{j=1}^{d^2} \left(1 - \frac{3p}{4}\right)^{2L-w[W_j]} \\ &\quad \times \left(\frac{p}{4}\right)^{w[W_j]} |\text{Tr}(W_j \Pi_{\mathcal{S}})|^2. \end{aligned} \quad (4.38)$$

A straightforward upper bound can be obtained by noting that the only operators with non-zero trace in $\mathcal{S}$ are tensor products of $X_p X_{p+L}$, $Y_p Y_{p+L}$, and $Z_p Z_{p+L}$. For these, one can upper bound $|\text{Tr}(W_j \Pi_{\mathcal{S}})|^2 \leq d_s^2$ (a severe overestimate) to find

$$\begin{aligned} \sum_{i=1}^{d^2} |\text{Tr}(K_i^{\mathcal{S}})|^2 &\leq d_s^2 \sum_{j=1}^{d^2} \left(1 - \frac{3p}{4}\right)^{2L-w[W_j]} \left(\frac{p}{4}\right)^{w[W_j]} \\ &= \sum_{w=0,2,\dots}^{2L} \left(1 - \frac{3p}{4}\right)^{2L-w} \left(\frac{p}{4}\right)^w \binom{L}{w/2} 3^{w/2} \\ &= d_s^2 \left[\sqrt{\left(1 - \frac{3p}{4}\right)^2 + 3 \left(\frac{p}{4}\right)^2} \right]^n. \end{aligned} \quad (4.39)$$

A lower bound is obtained by retaining only the $i = 1$ term in Eq. (4.10) of the main text,

$$\sum_{i=1}^{d^2} |\text{Tr}(K_i^{\mathcal{S}})|^2 \geq |\text{Tr}(\Pi_{\mathcal{S}})|^2 \left(1 - \frac{3p}{4}\right)^n = d_s^2 \left(1 - \frac{3p}{4}\right)^n, \quad (4.40)$$

which yields the bounds

$$\begin{aligned} \left(1 - \frac{3p}{4}\right)^n &\leq F(\mu_{\mathcal{S}}) \\ &\leq \frac{1}{1+d_s} + \frac{d_s}{1+d_s} \left[\sqrt{\left(1 - \frac{3p}{4}\right)^2 + 3\left(\frac{p}{4}\right)^2} \right]^n \end{aligned} \quad (4.41)$$

For small p , expanding the square root gives

$$\sqrt{\left(1 - \frac{3p}{4}\right)^2 + 3\left(\frac{p}{4}\right)^2} \approx 1 - \frac{3p}{4} + \frac{3p^2}{32} + \mathcal{O}(p^3), \quad (4.42)$$

so that in the large- n limit the upper bound becomes

$$\begin{aligned} &\frac{1}{1+d_s} + \frac{d_s}{1+d_s} \left(1 - \frac{3p}{4}\right)^n + \mathcal{O}(p^2) \\ &= \left(1 - \frac{3p}{4}\right)^n + \mathcal{O}(d_s^{-1}), \end{aligned} \quad (4.43)$$

and the two bounds coincide asymptotically. To leading order in p and for large n ,

$$F(\mu_{\mathcal{S}}) \approx \left(1 - \frac{3p}{4}\right)^n \approx F(\mu_{\mathcal{H}}), \quad (4.44)$$

where the second approximation holds because the upper bound on $F(\mu_{\mathcal{S}})$ exceeds $F(\mu_{\mathcal{H}})$ only by subleading corrections. Thus, at least asymptotically in n and for small p , the fidelity of a typical state is approximately equal to that of a scar subspace state, the latter being efficiently measurable. Away from this regime, however, particularly for larger p , the bound becomes less useful since the upper bound on $F(\mu_{\mathcal{S}})$ *exceeds* $F(\mu_{\mathcal{H}})$. Hence, we now construct bounds that are tight away from the small p regime.

4.C.2 Bound tight near $p = 1$

To extend our result beyond the large n and small p regime, we proceed by finding tighter lower and upper bounds for $|\text{Tr}(\Pi_{\mathcal{S}}W)|$, where W is a Pauli matrix, near

$p = 1$ (more precisely when $p \sim \mathcal{O}(1)$). First, note that this trace is non-vanishing only if W is a tensor product of the stabilizer generators defining the QMBS subspace, i.e., if $W \in \{\mathbb{I}, X_i X_{i+L}, Y_i Y_{i+L}, Z_i Z_{i+L}\}^{\otimes L}$ and for these the trace can be estimated. Specifically, we use the representation

$$W = \prod_{i \in S_x \subset [L]} X_i X_{i+L} \prod_{j \in S_z \subset [L]} Z_j Z_{j+L} \quad (4.45)$$

where $[L] := \{1, \dots, L\}$ denotes the set of indices of sites in the lattice (see Fig. 4.6). The scar subspace projector can be written in terms of the basis states in Eq. (4.12) of the main text,

$$\Pi_S = \sum_{\alpha=1}^4 \sum_{k=1}^L |\varphi_{\alpha,k}\rangle \langle \varphi_{\alpha,k}|, \quad (4.46)$$

so that

$$\text{Tr}(\Pi_S W) = \sum_{\alpha,k} \prod_{i \in S_x} \prod_{j \in S_z} \langle \varphi_{\alpha,k} | X_i X_{i+L} Z_j Z_{j+L} | \varphi_{\alpha,k} \rangle, \quad (4.47)$$

where we have used that if $|\psi\rangle$ is an eigenstate of operators A, B , then $\langle \psi | AB | \psi \rangle = \langle \psi | A | \psi \rangle \langle \psi | B | \psi \rangle$. To compute Eq. (4.47), the key step is to rewrite the sum over all basis states in the subspace as a sum over Pauli operators. In particular, note that the expectation value of an operator labeled by $1 \leq k \leq L$ can be mapped to the expectation value of a corresponding operator with respect to a reference state labeled $k = 1$, e.g.,

$$\begin{aligned} & \langle \varphi_{\alpha,k} | X_i X_{i+L} Z_j Z_{j+L} | \varphi_{\alpha,k} \rangle \\ &= \langle \varphi_{\alpha,1} | X_{i-k+1} X_{i+L-k+1} Z_{j-k+1} Z_{j+L-k+1} | \varphi_{\alpha,1} \rangle. \end{aligned} \quad (4.48)$$

Using the above equation and defining the cyclic permutation (in k) of an operator,

$$\mathbb{P}(W) \equiv \sum_{k'=0}^{L-1} \prod_{i \in S_x} X_{i-k'} X_{i+L-k'} \prod_{j \in S_z} Z_{j-k'} Z_{j+L-k'} , \quad (4.49)$$

which allows to write Eq. (4.47) as

$$\text{Tr}(\Pi_{\mathcal{S}}W) = \sum_{\alpha=1}^4 \langle \varphi_{\alpha} | \mathbb{P}(W) | \varphi_{\alpha} \rangle, \quad (4.50)$$

where, from now on, we abbreviate $|\varphi_{\alpha}\rangle \equiv |\varphi_{\alpha,1}\rangle$. Note that the following relations hold,

$$|\varphi_2\rangle = X_1 X_2 \dots X_L |\varphi_1\rangle \quad (4.51a)$$

$$|\varphi_3\rangle = X_2 \dots X_L |\varphi_1\rangle \quad (4.51b)$$

$$|\varphi_4\rangle = X_1 |\varphi_1\rangle \quad (4.51c)$$

which allow us to write

$$\begin{aligned} \text{Tr}(\Pi_{\mathcal{S}}W) = \langle \varphi_1 | [& X_1 \dots X_L \mathbb{P}(W) X_1 \dots X_L \\ & + X_2 \dots X_L \mathbb{P}(W) X_2 \dots X_L \\ & + X_1 \mathbb{P}(W) X_1 + \mathbb{P}(W)] | \varphi_1 \rangle. \end{aligned} \quad (4.52)$$

From this expression, we now show that the traces for different operators W can be readily estimated by distinguishing three cases: (i) W contains no $Z_i Z_{i+L}$ terms, (ii) W contains an odd number of $Z_i Z_{i+L}$ terms, and (iii) W contains an even number of $Z_i Z_{i+L}$ terms.

(i) W **contains no $Z_i Z_{i+L}$ terms.** In this case,

$$W = \prod_{i \in S_x} X_i X_{i+L} \quad (4.53)$$

which implies that

$$\begin{aligned} & X_1 \dots X_L \mathbb{P}(W) X_1 \dots X_L + X_2 \dots X_L \mathbb{P}(W) X_2 \dots X_L \\ & + X_1 \mathbb{P}(W) X_1 + \mathbb{P}(W) = 4\mathbb{P}(W), \end{aligned} \quad (4.54)$$

and thus

$$\text{Tr}(\Pi_{\mathcal{S}}W) = 4 \langle \varphi_1 | \mathbb{P}(W) | \varphi_1 \rangle. \quad (4.55)$$

To compute this note that $\langle \varphi_1 | X_i X_{i+L} | \varphi_1 \rangle = -1$ if $i \neq 1$ and 1 otherwise. Thus for W in Eq. (4.53),

$$\langle \varphi_1 | W | \varphi_1 \rangle = \begin{cases} (-1)^{|S_x|} & \text{if } i \neq 1, \\ -(-1)^{|S_x|} & \text{if } i = 1, \end{cases} \quad (4.56)$$

where $|S_x|$ is the cardinality of the set S_x . Since $\mathbb{P}(W)$ contains all permutations in k of W , it is a sum of L terms of which $|S_x|$ contain $X_1 X_{1+L}$ and $L - |S_x|$ do not. Consequently,

$$\begin{aligned} \text{Tr}(\Pi_S W) &= 4(L - |S_x|)(-1)^{|S_x|} - 4|S_x|(-1)^{|S_x|} \\ &= 4(-1)^{|S_x|}(L - 2|S_x|), \end{aligned} \quad (4.57)$$

for operators of the form Eq. (4.53).

(ii) **W contains an odd number of $Z_i Z_{i+L}$ terms.** In this case W is of the form Eq. (4.45), but $|S_z| = 2m + 1$ is odd. The terms in the cyclic permutation, $\mathbb{P}(W)$, now contain $Z_i Z_{i+L}$, and they can be grouped into two kinds: We label with $(1, \bar{1})$ any W which contain $Z_1 Z_{1+L}$ as well as an even number of $Z_i Z_{i+L}$ for $i \in \{2, \dots, L\}$. We label with $(\bar{1}, \bar{1})$ terms that contain no $Z_1 Z_{1+L}$ but an odd number of $Z_i Z_{i+L}$ for $i \in \{2, \dots, L\}$.

Using the fact that X and Z anti-commute, no matter the $X_i X_{i+L}$ content of W , we can tabulate the sign of the terms in Eq. (4.49) (to calculate Eq. (4.52)) for the two types below

Operator	Type $(1, \bar{1})$	Type $(\bar{1}, \bar{1})$
$\mathbb{P}(W)$	+	+
$X_1 \mathbb{P}(W) X_1$	−	+
$X_{1\dots L} \mathbb{P}(W) X_{1\dots L}$	−	−
$X_{2\dots L} \mathbb{P}(W) X_{2\dots L}$	+	−
	= 0	= 0

indicating that, for given W , these terms cancel in Eq. (4.52) and can therefore be neglected.

(iii) W **contains an even number of $Z_i Z_{i+L}$ terms.** In this case, a similar table can be constructed,

Operator	Type $(1, \bar{1})$	Type $(\bar{1}, \bar{1})$
$\mathbb{P}(W)$	+	+
$X_1 \mathbb{P}(W) X_1$	−	+
$X_{1\dots L} \mathbb{P}(W) X_{1\dots L}$	+	+
$X_{2\dots L} \mathbb{P}(W) X_{2\dots L}$	−	+

but now only terms of type $(1, \bar{1})$ are canceled by the symmetrization. Since Pauli expectation values are bounded in magnitude by 1, for operators containing an even number of $Z_i Z_{i+L}$ terms, the trace satisfies

$$\text{Tr}(\Pi_{\mathcal{S}} W) \leq 4(L - |S_z|). \quad (4.58)$$

To obtain a lower bound, we note that since L is always odd in our model, the total number of surviving Pauli terms is odd. As each such term contributes ± 1 , the sum cannot vanish, and therefore

$$1 \leq (\text{Tr}(\Pi_{\mathcal{S}} W))^2. \quad (4.59)$$

Indeed, in the extreme case all but one term cancels, leaving a net contribution of ± 1 .

Thus in summary, we get the exact expression

$$\text{Tr}(\Pi_{\mathcal{S}} W) = \begin{cases} 4(L - 2|S_x|)(-1)^{|S_x|} & \text{if } |S_z| = 0 \\ 0 & \text{if } |S_z| \text{ odd,} \end{cases} \quad (4.60)$$

and, if $|S_z| \neq 0$ and $|S_z|$ even, we obtain the following bounds, irrespective of S_x ,

$$1 \leq \text{Tr}(\Pi_{\mathcal{S}} W)^2 \leq 16(L - |S_z|)^2. \quad (4.61)$$

Thus Eq. (4.35) can be expressed,

$$\begin{aligned}
\sum_{i=1}^{d^2} \text{Tr}(K_i^s K_i^{\dagger s}) &= \frac{1}{d} \sum_{i=1}^{d^2} (1-p)^{w(W_i)} [\text{Tr}(\Pi_{\mathcal{S}} W_i)]^2 \\
&= \frac{1}{d} \sum_{S_x \subseteq [L]} \sum_{S_z \subseteq [L]} (1-p)^{w[W_{S_x, S_z}]} [\text{Tr}(\Pi_{\mathcal{S}} W_{S_x, S_z})]^2, \\
&= \frac{1}{d} \sum_{S_x \subseteq [L]} \sum_{\substack{S_z \subseteq [L] \\ |S_z|=\text{even}}} (1-p)^{2(|S_x|+|S_z|)} [\text{Tr}(\Pi_{\mathcal{S}} W_{S_x, S_z})]^2, \\
&= \frac{1}{d} \sum_{\substack{S_x \subseteq [L] \\ S_z = \emptyset}} (1-p)^{2|S_x|} [4(L-2|S_x|)]^2 \\
&\quad + \frac{1}{d} \sum_{\substack{S_x, S_z \subseteq [L] \\ |S_z| \neq 0, \\ |S_z|=\text{even}}} (1-p)^{2|S_x|+2|S_z|} \text{Tr}(W_{S_x S_z} \Pi_{\mathcal{S}})^2
\end{aligned} \tag{4.62}$$

where we have now written W_{S_x, S_z} to make the dependence clear. There are $\binom{L}{|S_x|}$ terms that contain $X_i X_{i+L}$ $|S_x|$ times, and similarly for S_z . Hence, we arrive at the following upper bound

$$\begin{aligned}
\sum_{i=1}^{d^2} \text{Tr}(K_i^s K_i^{\dagger s}) &\leq \frac{16}{d} \left[\sum_{|S_x|=0}^L (1-p)^{2|S_x|} \binom{L}{|S_x|} (L-2|S_x|)^2 \right. \\
&\quad \left. + [1 + (1-p)^2]^L \sum_{|S_z|=2, \text{even}}^{L-1} (1-p)^{2|S_z|} \binom{L}{|S_z|} (L-|S_z|)^2 \right]
\end{aligned} \tag{4.63}$$

where we used the binomial identity

$$\sum_{|S_x|=0}^L (1-p)^{2|S_x|} \binom{L}{|S_x|} = [1 + (1-p)^2]^L. \tag{4.64}$$

Eq. (4.63) can be evaluated using *Mathematica* yielding the closed form

$$\begin{aligned}
\sum_{i=1}^{d^2} \text{Tr}(K_i^s K_i^{\dagger s}) &\leq \frac{8}{d} L t^{2L-2} \left[(L+t-1) + (L-t+1)y^{L-2} \right. \\
&\quad \left. + \frac{2(L-1)(y^2-1)}{t^{L-2}} \right],
\end{aligned} \tag{4.65}$$

where $t = 1 + (1 - p)^2$ and $y = (2 - t)/t$. Similarly, the lower bound can be evaluated

$$\sum_{i=1}^{d^2} \text{Tr}(K_i^s K_i^{\dagger s}) \geq \frac{8}{d} t^{2L} \left[1 + y^L + \frac{2(L-1)(1+Ly^2)}{t^L} \right], \quad (4.66)$$

Because we only bound the traces of terms where $|S_z| \neq 0$ and $|S_z|$ even, the bounds derived above in Eqs. (4.65-4.66) are looser than the ones derived in the previous subsection in Eqs. (4.36-4.37) near $p = 0$. However, because we explicitly track the weights $(1 - p)^{w[W_i]}$, the bounds in Eqs. (4.65-4.66) become tighter for larger values of p . Since both sets of bounds hold for all $0 \leq p \leq 1$, we obtain uniformly tighter bounds by taking their intersection.

The term in Eq.(4.38) can be expressed as,

$$\begin{aligned} \sum_{i=1}^{d^2} |\text{Tr}(K_i^{\mathcal{S}})|^2 &= \sum_{i=1}^{d^2} \left(1 - \frac{3p}{4}\right)^{2L-w[W_i]} \left(\frac{p}{4}\right)^{w[W_i]} |\text{Tr}(W_i \Pi_{\mathcal{S}})|^2, \\ &= \sum_{\substack{S_x \subseteq [L] \\ S_z = \emptyset}} \left(1 - \frac{3p}{4}\right)^{2(L-|S_x|)} \left(\frac{p}{4}\right)^{2|S_x|} |\text{Tr}(\Pi_{\mathcal{S}} W_{S_x \emptyset})|^2 \\ &\quad + \sum_{S_x \subseteq [L]} \sum_{\substack{S_z \subseteq [L] \\ |S_z|=\text{even} \\ |S_z| \neq 0}} \left(1 - \frac{3p}{4}\right)^{2(L-|S_x|-|S_z|)} \left(\frac{p}{4}\right)^{2|S_x|+2|S_z|} \\ &\quad \times |\text{Tr}(\Pi_{\mathcal{S}} W_{S_x S_z})|^2. \end{aligned} \quad (4.67)$$

Using the same lines of reasoning as above, we arrive at the following upperbound,

$$\begin{aligned} \sum_{i=1}^{d^2} |\text{Tr}(K_i^{\mathcal{S}})|^2 &\leq 16 \sum_{|S_x|=0}^L \left(1 - \frac{3p}{4}\right)^{2(L-|S_x|)} \left(\frac{p}{4}\right)^{2|S_x|} \\ &\quad \times (L - 2|S_x|)^2 + 16 \left[\left(1 - \frac{3p}{4}\right)^2 + \left(\frac{p}{4}\right)^2 \right]^L \\ &\quad \times \left(\sum_{|S_z|=2, \text{even}}^L \binom{L}{|S_z|} \left(1 - \frac{3p}{4}\right)^{-2|S_z|} \left(\frac{p}{4}\right)^{2|S_z|} (L - |S_z|)^2 \right), \end{aligned} \quad (4.68)$$

which when evaluated using *Mathematica* gives,

$$\sum_{i=1}^{d^2} |\text{Tr}(K_i^S)|^2 \leq 8Lr^L \left[\frac{8(1-L)(1-w)}{w^2} + (L+w-1)w^{L-2} + (L+1-w)(2-w)^{L-2} \right], \quad (4.69)$$

where, $r = 1 - \frac{3p}{2} + \frac{5p^2}{8}$ and $w = 1 + (\frac{p}{4-3p})^2$. Similarly the lower bound can be evaluated as,

$$\sum_{i=1}^{d^2} |\text{Tr}(K_i^S)|^2 \geq 8Lr^L \left[w^{L+1} + (2-w)^L + \frac{8L(L-1)(1-w)}{w^2} \right]. \quad (4.70)$$

It is easy to verify that the upper and lower bounds on the second term derived in Eqs.(4.69) and (4.70) are tighter than the ones derived in Eqs.(4.39) and (4.40), for all $0 \leq p \leq 1$, respectively. Thus, putting everything together, the final bounds are

$$\begin{aligned} & \frac{1}{d_s(d_s+1)} [\max((4.37), (4.66)) + (4.70)] \leq F(\mu_S) \\ & \leq \frac{1}{d_s(d_s+1)} [\min((4.36), (4.65)) + (4.69)] . \end{aligned} \quad (4.71)$$

4.C.3 Concentration of Haar fidelity

In the main text, we analyze the average fidelity. Here we provide additional arguments that this estimate holds not only on average but for the overwhelming majority of states. Whether this alone is sufficient for quantum simulation is a separate question, as quantum-simulated states may constitute a highly structured subset of Hilbert space. For the present application to thermalization, however, this assumption is well motivated. In any case, our numerical results provide strong evidence that the fidelity equivalence also holds for specific quantum-simulated states.

Levy's lemma [253], implies that the probability of a state having channel a fidelity that is far from the average channel fidelity is exponentially small. Specifically, for a

system with n qubits,

$$\begin{aligned} & \Pr_{|\phi\rangle \sim \mu_H} [|\langle \phi | \mathcal{E}(|\phi\rangle\langle\phi|) | \phi\rangle - \mathbb{E}_{|\psi\rangle \sim \mu_H} [\langle \psi | \mathcal{E}(|\psi\rangle\langle\psi|) | \psi\rangle] | \geq \varepsilon] \\ & \leq 2 \exp\left(-\frac{2^n \varepsilon^2}{72\pi^3}\right) \end{aligned} \quad (4.72)$$

The use of Levy's lemma requires that the channel fidelity is a Lipschitz continuous function. While this is a well-known fact, we provide an elementary proof for the sake of completeness.

Lemma 4.1 *The function $f : \mathbb{C}^d \rightarrow \mathbb{R}$*

$$f(|\psi\rangle) := \langle \psi | \mathcal{E}(|\psi\rangle\langle\psi|) | \psi \rangle, \quad (4.73)$$

is Lipschitz continuous, with,

$$|f(|\psi\rangle) - f(|\phi\rangle)| \leq 4 \| |\psi\rangle - |\phi\rangle \|_2 \quad (4.74)$$

Proof. To prove the above, we first add and subtract the cross term $\text{Tr} [|\phi\rangle\langle\phi| \Phi(|\psi\rangle\langle\psi|)]$ to bring the expression to a nice form. Then, we use matrix Hölder inequality, $\text{Tr}(A^\dagger B) \leq \|A\|_\infty \|B\|_1$ to derive an upper bound in terms of Schatten-1 norm (or trace distance),

$$\begin{aligned} & |f(|\psi\rangle) - f(|\phi\rangle)| \\ &= \text{Tr} [|\psi\rangle\langle\psi| \Phi(|\psi\rangle\langle\psi|)] - \text{Tr} [|\phi\rangle\langle\phi| \Phi(|\phi\rangle\langle\phi|)], \\ &= \text{Tr} [(|\psi\rangle\langle\psi| - |\phi\rangle\langle\phi|) \Phi(|\psi\rangle\langle\psi|)] \\ &\quad - \text{Tr} [|\phi\rangle\langle\phi| \Phi(|\psi\rangle\langle\psi| - |\phi\rangle\langle\phi|)], \\ &\leq \| |\psi\rangle\langle\psi| - |\phi\rangle\langle\phi| \|_1 + \|\Phi(|\psi\rangle\langle\psi| - |\phi\rangle\langle\phi|)\|_1, \\ &\leq 2 \| |\psi\rangle\langle\psi| - |\phi\rangle\langle\phi| \|_1, \\ &\leq 4 \| |\psi\rangle - |\phi\rangle \|_2. \end{aligned} \quad (4.75)$$

In the 5-th line, we use the data processing inequality, by noting that the trace distance $2T(\rho, \sigma) = \|\rho - \sigma\|_1$. Finally, in the 6-th line we have used the arguments in Ref. [252] (section 9.2, Example 54), that show that $\| |\psi\rangle\langle\psi| - |\phi\rangle\langle\phi| \|_1 \leq 2 \| |\psi\rangle - |\phi\rangle \|_2$.

4.D MCMC Details

In this Appendix, we describe the Markov Chain Monte Carlo (MCMC) sampling procedure used for DFE, Eq. (4.1), of arbitrary states in the stabilizer scar subspace. Ref. [110] guarantees that, if the relevant expectation values are known exactly, the number of MCMC samples required is a constant that depends only on the desired precision ϵ and failure probability δ . If these expectation values are instead estimated experimentally from a finite number of measurement shots, then, following Ref. [112] and the arguments given in the main text, the number of copies of the quantum state required additionally scales polynomially with the number of qubits, i.e., the protocol is efficient.

While Refs. [110, 112] establish theoretical guarantees for the DFE procedure, they assume access to i.i.d. samples from the target state Pauli distribution $P_\rho(k)$. Generating such samples is not trivial since there are exponentially many Paulis to sample from. In this manuscript, we use a Metropolis-Hastings (MH) algorithm that proposes Pauli operators with non-zero support on $\mathcal{S}$ with *uniform probability*, clearly not an efficient proposal scheme. Alternative schemes will be presented in a forthcoming publication.

4.D.1 Metropolis-Hasting algorithm

The algorithm proposes only Pauli operators that have non-zero matrix elements within the target scar subspace, i.e., operators drawn from the set

$$P_{\mathcal{S}} := \{W_k \in \mathcal{P}_n \mid \exists \rho \in \mathcal{S} \text{ with } \text{Tr}[\rho W_k] \neq 0\} \quad (4.76)$$

For the prototype model of Eqs. (4.11-4.12), as we show in Eq. (4.79) below, this set contains $|P_{\mathcal{S}}| = (n^2 - n + 2)d$ operators, still exponentially many in n , but a small fraction of the full set of d^2 Pauli operators.

To ensure that only Pauli operators from the set $P_{\mathcal{S}}$ are proposed, we use a decision

tree, which is summarized compactly in Fig. 4.7. We seek Pauli operators with non-zero matrix elements between pairs of scar basis states. It is useful to characterize each basis state in terms of a “defect” located at site $k = 1, \dots, L$ with one of four flavors $a = 1, 2, 3, 4$, corresponding to the four panels of Fig. 4.6(b). The allowed Pauli operators then either move the defect and/or change its flavor, act as a diagonal operator. They are parameterized as

$$W = (-i)^{\mathbf{a} \cdot \mathbf{b}} X^{a_1} \dots X^{a_{2L}} Z^{b_1} \dots Z^{b_{2L}}, \quad (4.77)$$

where $\mathbf{a}, \mathbf{b} \in \{0, 1\}^{2L}$. In the first step, Fig. 4.7(i), the pattern of Z operators is determined: with probability $p = \frac{2+2L}{2+2L+4L(L-1)}$ no “single” Z operators are generated, while with probability $1 - p$ two “single” Z operators are placed at two randomly chosen, distinct locations $i \in \{1, \dots, L\}$ (the “upper row” of the $L \times 2$ lattice in Fig. 4.6(b)), corresponding to a “hopping” of a defect. The operator W is then “dressed” by multiplication with a random product of $\{\mathbb{I}_i \mathbb{I}_{i+L}, Z_i Z_{i+L}\}^{\otimes L}$, where the number $k \in \{0, \dots, L\}$ of $Z_i Z_{i+L}$ stabilizer factors included is drawn from a binomial distribution, i.e., there are $\binom{L}{k}$ ways of placing k such stabilizer operators, see Fig. 4.7(ii).

This is followed by the placement of X operators in Fig. 4.7(iii): If no single Z operator was placed in the first step (right column), then the defect does not hop, but its flavor may still change. In this case, one of four X patterns is chosen:

$$\mathbb{I} \quad \text{with probability } \frac{1}{2L+2}, \quad (4.78a)$$

$$\prod_{i=1}^L X_i \quad \text{with probability } \frac{1}{2L+2}, \quad (4.78b)$$

$$X_k \quad \text{with probability } \frac{L}{2L+2}, \quad (4.78c)$$

$$\prod_{\substack{i=1 \\ i \neq k}}^L X_i \quad \text{with probability } \frac{L}{2L+2}, \quad (4.78d)$$

where, in the last two cases, the location $k = 1, \dots, L$ is chosen with equal probability.




Figure 4.7: *MCMC suggest-step decision tree*. Shown is the decision tree used to uniformly generate Pauli operators with nonzero expectation value within the stabilizer scar subspace. This proceeds in four steps: (i) a Z operator pattern is generated, in which either a “defect hops” (left column) or does not (right column); (ii) this Z pattern is further dressed by a random product of $\{I_i I_{i+L}, Z_i Z_{i+L}\}^{\otimes L}$ operators; (iii) based on the resulting Z pattern, an X pattern is generated that either connects the defect’s origin and target (left column, not shown are the inverted X patterns in the top row) or merely changes its flavor (right column); and (iv) this X pattern is similarly dressed by a random product of $\{I_i I_{i+L}, X_i X_{i+L}\}^{\otimes L}$ operators. The corresponding matrix elements, together with the phase obtained from Eq. (4.77), are tracked throughout to construct a $4L \times 4L$ matrix in the QMBS subspace.

If, on the other hand, two “single” Z operators were placed in the first step (left column), the X pattern is largely fixed: it must accommodate the “hopping” of the defect, although the flavor may still change in this process. In general, this requires placing an X string connecting the two “single” Z locations, which correspond to the origin and target of the defect hop. Owing to the periodic boundary conditions, this can be done in two ways; in addition, the X string may or may not include either of the two “single” Z locations, with each choice corresponding to a different combination of origin/target flavor change. This yields 8 possibilities in total. Finally in Fig. 4.7(iv), as before, the resulting operator is “dressed” by multiplication with a random product of $\{\mathbb{I}_i \mathbb{I}_{i+L}, X_i X_{i+L}\}^{\otimes L}$. As the decision tree is traversed, the matrix elements between (sets of) basis states are tracked, the phase appearing in Eq. (4.77) is computed, and a corresponding $4L \times 4L$ matrix is constructed, corresponding to the projection of W into the scar subspace.

This prescription allows one to count the number of nonzero Pauli observables supported on the subspace by summing over all branches of the decision tree in Fig. 4.7. There are $\binom{L}{2} = L(L-1)/2$ ways of placing the single- Z operators, each admitting 8 corresponding X -operator patterns (left column of the figure), together with $2 + 2L$ configurations containing only flavor-changing X patterns. Each such pattern can further be dressed independently by 2^L products from $\{\mathbb{I}_i \mathbb{I}_{i+L}, X_i X_{i+L}\}^{\otimes L}$ and 2^L products from $\{\mathbb{I}_i \mathbb{I}_{i+L}, Z_i Z_{i+L}\}^{\otimes L}$, yielding

$$|P_S| = 2^{2L} \left(8 \binom{L}{2} + 2 + 2L \right) = d(n^2 - n + 2). \quad (4.79)$$

The proposed t -th sample $k_{t+1}^{\text{prop.}}$ is accepted or rejected based on the ratio [242]

$$\alpha = \frac{P_\rho(k_{t+1}^{\text{prop.}})}{P_\rho(k_t)}, \quad P_\rho(k) = \chi_\rho^2(k) = \frac{\text{Tr}[\rho P_k]^2}{d} \quad (4.80)$$

where k_t is the current sample in the MCMC. Drawing a uniform random number $u \in [0, 1]$, the proposal is accepted, $k_{t+1} = k_{t+1}^{\text{prop.}}$, if $\alpha > u$, and rejected, $k_{t+1} = k_t$, otherwise (i.e., if $\alpha \leq u$).

4.D.2 I.i.d. MCMC samples

As is inherent to MCMC methods, successive samples are generally autocorrelated because each sample is generated from the previous one [254]. To quantify this effect, we compute the Pearson correlation coefficient of the sampled estimators $E_\tau \equiv \chi_\sigma(k_\tau)/\chi_\rho(k_\tau)$. Specifically, the autocorrelation at lag τ is defined as the ratio of the covariance between samples separated by τ to the variance of the series,

$$C_\tau \equiv \frac{\text{Cov}(E_t, E_{t+\tau})}{\text{Var}(E_t)} = \frac{\sum_{t=1}^{N-\tau} (E_t - \bar{E})(E_{t+\tau} - \bar{E})}{\sum_{t=1}^N (E_t - \bar{E})^2}, \quad (4.81)$$

where $\bar{E}$ is the sample mean and N is the total number of samples in the chain. The autocorrelation determines the effective sample size N_{eff} , which is generally smaller than the total number of samples N , thereby determining how fast the sample mean converges to the true value. For our numerical simulations we discard the first 1000 samples as a burn-in period and thin the chain, keeping only every 100-th sample thereafter. We find that the auto-correlation factor $C_{100} \approx 0$, indicating a negligible residual autocorrelation. Consequently, we treat our samples as being generated independently, i.e., $N_{\text{eff}} \approx N$.

When the autocorrelation cannot be made negligible through thinning, the effective sample size can be estimated analytically. Concretely, consider the random variable $\bar{E} = \frac{1}{N} \sum_i E_i$, where E_i is the random variable associated with the MCMC sampling, then

$$\text{Var}(\bar{E}) = \frac{\text{Var}(E_t)}{N^2} \left(N + 2 \sum_{t=1}^N \sum_{\tau=1}^{N-t} C_\tau \right), \quad (4.82)$$

which recovers the familiar $1/N$ scaling if the second term vanishes. One can write

$$\sum_{t=1}^N \sum_{\tau=1}^{N-t} C_\tau = \sum_{\tau=1}^N \sum_{t=1}^{N-\tau} C_\tau = \sum_{\tau=1}^N (N - \tau) C_\tau, \quad (4.83)$$

and, parameterizing $C_\tau \equiv \epsilon_{\text{acf}} e^{-k(\tau-1)}$ where $\epsilon_{\text{acf}} = C_1$, and k a positive constant

which can be extracted from numerics, one gets

$$\text{Var}(\bar{E}) = \frac{\text{Var}(E_t)}{N} \left(1 + \frac{2\epsilon_{\text{acf}}}{N} \sum_{\tau=1}^N (N - \tau) e^{-k\tau} \right). \quad (4.84)$$

4.D.3 Simulating an experiment with shots

To demonstrate that the number of shots required to achieve the desired precision falls below the bound in Eq. (4.7) we classically emulate the shot-based measurements of expectation values for a test state σ and target state $\rho \in \mathcal{S}$. Recall that ρ is the ideal state for which expectation values are known exactly, while σ is the state that is estimated in experiment with a-priori unknown expectation values. The following analysis requires knowledge of the exact expectation value $\text{Tr}(W_i\sigma)$, which is inaccessible in practice; it can be understood, in a sense, as the ideal (i.e. lower bound) number of shots, while the results in Eq. (4.7) of the main text give a worse-case guarantee (upper bound). Given the expectation value $\langle W_{k_i} \rangle_\sigma \equiv \text{Tr}[\sigma W_{k_i}]$, we define the outcome probabilities

$$p_{\pm}^{(i)} = \frac{1 \pm \langle W_{k_i} \rangle_\sigma}{2}, \quad (4.85)$$

which satisfy $p_+^{(i)} + p_-^{(i)} = 1$. Sampling from this Bernoulli distribution yields measurement outcomes corresponding to projective measurements of W_{k_i} .

To determine the number of shots m_i required to achieve a prescribed precision and confidence, we demand

$$\Pr(|\hat{Y} - Y| \geq \epsilon) \leq \delta, \quad (4.86)$$

where $Y = \frac{1}{s} \sum_{i=1}^s X_i$ and $\hat{Y} = \frac{1}{s} \sum_{i=1}^s \hat{X}_i$ are the estimators built from exact and shot-sampled expectation values, respectively, and $X_i, \hat{X}_i$ are defined below Eq. (4.1). Suppose that for each Pauli sample i we have $\Pr(|\hat{X}_i - X_i| \geq \epsilon_i) \leq \delta_i$. Then, by the triangle inequality and independence,

$$\Pr\left(\frac{1}{s} \sum_{i=1}^s |\hat{X}_i - X_i| \leq \frac{1}{s} \sum_{i=1}^s \epsilon_i\right) \geq \prod_{i=1}^s (1 - \delta_i). \quad (4.87)$$

Hence, Eq. (4.86) is satisfied provided each term obeys

$$\Pr\left(|\hat{X}_i - X_i| \geq \epsilon\right) \leq 1 - (1 - \delta)^{1/s}. \quad (4.88)$$

Following Ref. [110], the shot-based estimator reads

$$\hat{X}_i = \frac{1}{m_i \langle W_{k_i} \rangle_\rho} \sum_{j=1}^{m_i} A_{ij}, \quad (4.89)$$

where $A_{ij} \in \{+1, -1\}$ are i.i.d. draws from the distribution in Eq. (4.85) and $\langle W_{k_i} \rangle_\rho \equiv \text{Tr}(W_{k_i} \rho)$. The failure probability for sample i is therefore

$$\begin{aligned} \Pr\left(|\hat{X}_i - X_i| \geq \epsilon\right) &= \Pr\left(\left|\langle W_{k_i} \rangle_\rho - \frac{1}{m_i} \sum_{j=1}^{m_i} A_{ij}\right| \geq \epsilon \langle W_{k_i} \rangle_\rho\right) \\ &= F\left(\left\lfloor m_i \left(p_+^{(i)} - \frac{\epsilon}{2} \langle W_{k_i} \rangle_\rho\right) \right\rfloor; m_i, p_+^{(i)}\right) \\ &\quad + 1 - F\left(\left\lceil m_i \left(p_+^{(i)} + \frac{\epsilon}{2} \langle W_{k_i} \rangle_\rho\right) \right\rceil; m_i, p_+^{(i)}\right), \end{aligned} \quad (4.90)$$

where the binomial cumulative distribution function is,

$$F(a; m, p) = \sum_{n=0}^a \binom{m}{n} p^n (1-p)^{m-n}. \quad (4.91)$$

While the expression in Eq. (4.90) is not strictly monotonic in n , it is monotonic in aggregate. Hence, the required number of shots m_i for each expectation value $\langle W_{k_i} \rangle$ to achieve precision ϵ and failure probability δ is then obtained by finding the smallest m_i such that moving average of Eq. (4.90) does not exceed the right-hand side of Eq. (4.88).

Additional numerical studies complementing Fig. 4.5(b) of the main text are presented in Fig. 4.8 ($n = 10$ qubits, $\ell = 15$ layers, $p = 10^{-3}$). In (a), eight *different* randomly selected MCMC chains are shown, exhibiting behavior consistent with the results presented in the main text. For some chains, an apparent systematic deviation is observed at small s between the estimates obtained from exact Pauli expectation values (solid lines) and those obtained from shot-sampled measurements (circles).




Figure 4.8: *Additional noisy circuit simulation with shots.* (a) Results for $n = 10$ qubits, $\ell = 15$ circuit layers, and depolarizing noise strength $p = 0.001$, analogous to Fig. 4.5(b) of the main text. Eight different randomly selected MCMC chains are shown out of 100 generated. Circles denote shot-sampled estimates (10^4 shots per Pauli observable), while solid lines show the corresponding exact Pauli expectation values. Some shot-sampled trajectories exhibit an apparent systematic offset from the exact curves. This arises from the construction of the statistical uncertainty: the error bars represent the accumulated uncertainty of the fidelity estimate after s MCMC samples, obtained via error propagation. As a result, an early (or intermediate) statistical outlier in a Pauli expectation value can bias the cumulative estimate until its effect is gradually averaged out by subsequent samples. (b) Same as the bottom panel of Fig. 4.5(b) of the main text, showing the final fidelity precision after $s = 10^4$ MCMC samples.




Figure 4.9: *Shot dependence.* The same MCMC chain as in Fig. 4.5(b) of the main text applied to noisy circuit simulations with $n = 10$ qubits, $\ell = 15$ layers, and depolarizing noise strength $p = 10^{-3}$. The number of shots per Pauli observable is varied between 10^3 (10% of the shots used in the main text, left panel), $5 \cdot 10^3$ (50%, middle panel), and $5 \cdot 10^4$ (500%, right panel). The bottom panels show the resulting precision after $s = 10^4$ MCMC samples.

This behavior originates from the way the statistical uncertainty is computed: the shown error bars correspond to the accumulated uncertainty of the fidelity estimate after s MCMC samples, represented using the sample variance obtained via error propagation, and uncertainties at successive data points are not independent. Consequently, an atypical shot-based estimate, i.e., an outlier, at an early stage of the sampling process can bias the cumulative estimate for a number of subsequent MCMC samples, with the resulting deviation persisting until it is averaged out by additional samples. Ultimately, all trajectories converge, achieving the improved precision and failure probability that we observe is significantly better than the bounds derived in the main text.

Fig. 4.9 shows the same MCMC chains as in Fig. 4.5(b) of the main text. However, we vary the shots used to evaluate each Pauli from 10^3 (10% of the shots used in the main text, left panel), $5 \cdot 10^3$ (50%, middle panel), and $5 \cdot 10^4$ (500%, right panel). These results are also consistent with the bounds presented in the main text.

Chapter 5

CONCLUSION AND OUTLOOK

The work in this thesis was motivated by the current stage of quantum computing, in which programmable devices have grown capable enough to begin addressing physical systems in nuclear and high-energy physics. Specifically, it aimed to address two central obstacles that stand in the way of turning that prospect into reliable physics: the cost of executing the required circuits, and the difficulty of verifying their output in the regime where, by construction, no classical answer exists for comparison. The guiding principle throughout has been to exploit structure, whether physical knowledge available before a computation begins or special properties of particular quantum states, so that simulation becomes at once cheaper to perform and possible to verify.

The three main results chapters developed this idea in distinct ways. Chapter 2 exploited known structure: by recasting the Gaussian integral transform of a response function in a Fourier basis, whose moments are real-time expectation values, and by using the mean and variance of the energy distribution to center and rescale the reconstruction, the cost of resolving spectral densities was reduced by up to two orders of magnitude for realistic nuclear response functions. Chapter 3 turned from using structure to discovering it, identifying an exactly solvable family of quantum many-body scars in a $2 + 1$ -dimensional $\mathbb{Z}_2$ lattice gauge theory whose members are stabilizer states, carrying exactly zero magic and zero distillable entanglement while embedded in otherwise thermalizing dynamics, with their stability traced to a commutant algebra. Chapter 4 then put that discovered structure to work, using the exact classical solvability of the scar subspace and its vanishing non-stabilizerness to build a self-contained and scalable benchmark whose measured fidelity bounds the

fidelity of the intractable, non-scar dynamics running on the same device.

Taken together, these results work towards the promise of quantum simulation and its practical use from two angles: cost and verification. On the side of cost, exploiting prior physical knowledge showed that the resources demanded by a physically relevant calculation can be cut by orders of magnitude, bringing it closer to what near-term hardware can deliver. On the side of verification, the discovery of stabilizer scars supplied a classically solvable reference that lives inside an otherwise intractable computation, turning benchmarking from an apparent impossibility into a concrete and scalable procedure. Underlying both is a single lesson: that structure, whether known in advance or discovered within a system, can make quantum simulation not only cheaper but also trustworthy.

The scars of Chapter 3 suggest several directions for further work. Extending these constructions beyond quasi-one-dimensional $\mathbb{Z}_2$ settings to genuinely two-dimensional systems, other gauge groups, and non-Abelian theories is a natural next step. The non-Abelian case is especially interesting, because distillable entanglement is generically nonzero there even at zero coupling [45], so any scars would carry a different resource-theoretic fingerprint and could not be classically simple in the same exact sense. Since this work has been published, additional families of stabilizer scars have been discovered [103, 102], but a systematic understanding of which commutant algebras support stabilizer scars would clarify how special the present example is.

The benchmarking strategy of Chapter 4 also raises further questions. The present approach certifies a device on the scar subspace and extrapolates to the bulk under assumptions about how noise acts across sectors, so the tightness and the regime of validity of the scar-to-bulk fidelity bound deserve closer scrutiny than a single model can provide. Moreover, the bound also presumes a noise model that can be characterized, which is appropriate for near-term devices, but less clear in fault-tolerant settings, where logical errors are shaped by the code and the decoder, in ways that are not known a priori [255, 256]. Extending the embedded-reference scheme to

certify logical rather than physical dynamics, where the effective noise on the logical subspace is itself an open question, is a natural and important next step.

BIBLIOGRAPHY

- [1] Ulrich Schollwöck. “The density-matrix renormalization group in the age of matrix product states”. In: *Annals of Physics* 326.1 (2011). January 2011 Special Issue, pp. 96–192. ISSN: 0003-4916. DOI: <https://doi.org/10.1016/j.aop.2010.09.012>. URL: <https://www.sciencedirect.com/science/article/pii/S0003491610001752>.
- [2] Nicholas Metropolis et al. *Equation of state calculations by fast computing machines*. Tech. rep. Los Alamos Scientific Lab., Los Alamos, NM (United States); Univ. of Chicago, IL (United States), Mar. 1953. DOI: [10.2172/4390578](https://www.osti.gov/biblio/4390578). URL: <https://www.osti.gov/biblio/4390578>.
- [3] Christof Gattringer and Christian B. Lang. *Quantum Chromodynamics on the Lattice: An Introductory Presentation*. Vol. 788. Lecture Notes in Physics. Berlin, Heidelberg: Springer, 2010. ISBN: 978-3-642-01849-7. DOI: [10.1007/978-3-642-01850-3](https://doi.org/10.1007/978-3-642-01850-3).
- [4] S. Dürr et al. “Ab Initio Determination of Light Hadron Masses”. In: *Science* 322.5905 (2008), pp. 1224–1227. DOI: [10.1126/science.1163233](https://doi.org/10.1126/science.1163233). eprint: <https://www.science.org/doi/pdf/10.1126/science.1163233>. URL: <https://www.science.org/doi/abs/10.1126/science.1163233>.
- [5] Gert Aarts. “Introductory lectures on lattice QCD at nonzero baryon number”. In: *J. Phys. Conf. Ser.* 706.2 (2016), p. 022004. DOI: [10.1088/1742-6596/706/2/022004](https://doi.org/10.1088/1742-6596/706/2/022004). arXiv: 1512.05145 [hep-lat].
- [6] Masayuki Asakawa, Tetsuo Hatsuda, and Yasuyuki Nakahara. “Maximum entropy analysis of the spectral functions in lattice QCD”. In: *Prog. Part. Nucl.*

- Phys.* 46 (2001), pp. 459–508. DOI: 10.1016/S0146-6410(01)00150-8. arXiv: hep-lat/0011040 [hep-lat]. URL: <https://arxiv.org/abs/hep-lat/0011040>.
- [7] Matthias Troyer and Uwe-Jens Wiese. “Computational Complexity and Fundamental Limitations to Fermionic Quantum Monte Carlo Simulations”. In: *Phys. Rev. Lett.* 94 (17 May 2005), p. 170201. DOI: 10.1103/PhysRevLett.94.170201. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.94.170201>.
- [8] Richard P Feynman. “Simulating physics with computers”. In: *International Journal of Theoretical Physics* 21.6-7 (1982), pp. 467–488. DOI: 10.1007/BF02650179.
- [9] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information*. 10th Anniversary. Cambridge: Cambridge University Press, 2010. ISBN: 9781107002173.
- [10] Seth Lloyd. “Universal quantum simulators”. In: *Science* 273.5278 (1996), pp. 1073–1078.
- [11] Natalie Klco, Alessandro Roggero, and Martin J Savage. “Standard model physics and the digital quantum revolution: thoughts about the interface”. In: *Reports on Progress in Physics* 85.6 (May 2022), p. 064301. DOI: 10.1088/1361-6633/ac58a4. URL: <https://dx.doi.org/10.1088/1361-6633/ac58a4>.
- [12] Christian W Bauer et al. “Quantum simulation for high-energy physics”. In: *PRX quantum* 4.2 (2023), p. 027001.
- [13] Douglas Beck et al. “Quantum information science and technology for nuclear physics. Input into US Long-Range Planning, 2023”. In: *arXiv preprint arXiv:2303.00113* (2023).

- [14] M. C. Bañuls et al. “Simulating Lattice Gauge Theories within Quantum Technologies”. In: *Eur. Phys. J. D* 74.8 (2020), p. 165. DOI: 10.1140/epjd/e2020-100571-8. arXiv: 1911.00003 [quant-ph].
- [15] Alessandro Roggero and Joseph Carlson. “Dynamic linear response quantum algorithm”. In: *Phys. Rev. C* 100 (3 Sept. 2019), p. 034610. DOI: 10.1103/PhysRevC.100.034610. URL: <https://link.aps.org/doi/10.1103/PhysRevC.100.034610>.
- [16] Alessandro Roggero et al. “Preparation of excited states for nuclear dynamics on a quantum computer”. In: *Phys. Rev. C* 102 (6 Dec. 2020), p. 064624. DOI: 10.1103/PhysRevC.102.064624. URL: <https://link.aps.org/doi/10.1103/PhysRevC.102.064624>.
- [17] P. Krantz et al. “A quantum engineer’s guide to superconducting qubits”. In: *Applied Physics Reviews* 6.2 (2019), p. 021318.
- [18] Morten Kjaergaard et al. “Superconducting qubits: Current state of play”. In: *Annual Review of Condensed Matter Physics* 11 (2020), p. 369.
- [19] Juan I Cirac and Peter Zoller. “Quantum computations with cold trapped ions”. In: *Physical review letters* 74.20 (1995), p. 4091.
- [20] Colin D. Bruzewicz et al. “Trapped-ion quantum computing: Progress and challenges”. In: *Applied Physics Reviews* 6.2 (2019), p. 021314.
- [21] Antoine Browaeys and Thierry Lahaye. “Many-body physics with individually controlled Rydberg atoms”. In: *Nature Physics* 16.2 (2020), pp. 132–142.
- [22] Ehud Altman et al. “Quantum simulators: Architectures and opportunities”. In: *PRX quantum* 2.1 (2021), p. 017003.
- [23] M. Saffman, T. G. Walker, and K. Mølmer. “Quantum information with Rydberg atoms”. In: *Rev. Mod. Phys.* 82 (2010), p. 2313.

- [24] Peter W. Shor. “Scheme for reducing decoherence in quantum computer memory”. In: *Phys. Rev. A* 52 (4 Oct. 1995), R2493(R)–R2496(R). DOI: 10.1103/PhysRevA.52.R2493. URL: <https://link.aps.org/doi/10.1103/PhysRevA.52.R2493>.
- [25] Austin G. Fowler et al. “Surface codes: Towards practical large-scale quantum computation”. In: *Physical Review A* 86 (3 Sept. 2012), p. 032324. DOI: 10.1103/PhysRevA.86.032324. URL: <https://link.aps.org/doi/10.1103/PhysRevA.86.032324>.
- [26] Yosi Atia and Dorit Aharonov. “Fast-forwarding of Hamiltonians and exponentially precise measurements”. In: *Nature Communications* 8.1 (2017), p. 1572. DOI: 10.1038/s41467-017-01637-7. URL: <https://doi.org/10.1038/s41467-017-01637-7>.
- [27] Dorit Aharonov et al. “A Polynomial-Time Classical Algorithm for Noisy Random Circuit Sampling”. In: *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*. STOC 2023. Orlando, FL, USA: Association for Computing Machinery, 2023, pp. 945–957. ISBN: 9781450399135. DOI: 10.1145/3564246.3585234. URL: <https://doi.org/10.1145/3564246.3585234>.
- [28] John Preskill. “Quantum computing and the entanglement frontier”. In: *arXiv* (2012). arXiv: 1203.5813 [quant-ph].
- [29] Aram W Harrow and Ashley Montanaro. “Quantum computational supremacy”. In: *Nature* 549.7671 (2017), pp. 203–209.
- [30] Frank Arute et al. “Quantum supremacy using a programmable superconducting processor”. In: *Nature* 574.7779 (2019), pp. 505–510.
- [31] Dominik Hangleiter et al. “Direct certification of a class of quantum simulations”. In: *Quantum Science and Technology* 2.1 (2017), p. 015004.

- [32] Jose Carrasco et al. “Theoretical and experimental perspectives of quantum verification”. In: *PRX Quantum* 2.1 (2021), p. 010102.
- [33] Masuo Suzuki. “General theory of fractal path integrals with applications to many-body theories and statistical physics”. In: *Journal of Mathematical Physics* 32.2 (Feb. 1991), pp. 400–407. ISSN: 0022-2488. DOI: 10.1063/1.529425. eprint: https://pubs.aip.org/aip/jmp/article-pdf/32/2/400/19166143/400_1_online.pdf. URL: <https://doi.org/10.1063/1.529425>.
- [34] Guang Hao Low and Isaac L. Chuang. “Optimal Hamiltonian Simulation by Quantum Signal Processing”. In: *Phys. Rev. Lett.* 118 (1 Jan. 2017), p. 010501. DOI: 10.1103/PhysRevLett.118.010501. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.118.010501>.
- [35] Andrew M. Childs et al. “Toward the first quantum simulation with quantum speedup”. In: *Proceedings of the National Academy of Sciences* 115.38 (2018), pp. 9456–9461. DOI: 10.1073/pnas.1801723115. eprint: <https://www.pnas.org/doi/pdf/10.1073/pnas.1801723115>. URL: <https://www.pnas.org/doi/abs/10.1073/pnas.1801723115>.
- [36] Sonia Bacca and Saori Pastore. “Electromagnetic reactions on light nuclei”. In: *Journal of Physics G: Nuclear and Particle Physics* 41.12 (Nov. 2014), p. 123002. DOI: 10.1088/0954-3899/41/12/123002. URL: <https://doi.org/10.1088/0954-3899/41/12/123002>.
- [37] Victor D. Efros, Winfred Leidemann, and Giuseppina Orlandini. “Response functions from integral transforms with a Lorentz kernel”. In: *Physics Letters B* 338.2 (1994), pp. 130–133. ISSN: 0370-2693. DOI: [https://doi.org/10.1016/0370-2693\(94\)91355-2](https://doi.org/10.1016/0370-2693(94)91355-2). URL: <https://www.sciencedirect.com/science/article/pii/0370269394913552>.

- [38] Alessandro Roggero et al. “Quantum computing for neutrino-nucleus scattering”. In: *Phys. Rev. D* 101 (7 Apr. 2020), p. 074038. DOI: 10.1103/PhysRevD.101.074038. URL: <https://link.aps.org/doi/10.1103/PhysRevD.101.074038>.
- [39] John Kogut and Leonard Susskind. “Hamiltonian formulation of Wilson’s lattice gauge theories”. In: *Phys. Rev. D* 11 (2 Jan. 1975), pp. 395–408. DOI: 10.1103/PhysRevD.11.395. URL: <https://link.aps.org/doi/10.1103/PhysRevD.11.395>.
- [40] John B. Kogut. “An introduction to lattice gauge theory and spin systems”. In: *Rev. Mod. Phys.* 51 (4 Oct. 1979), pp. 659–713. DOI: 10.1103/RevModPhys.51.659. URL: <https://link.aps.org/doi/10.1103/RevModPhys.51.659>.
- [41] A Yu Kitaev. “Fault-tolerant quantum computation by anyons”. In: *Annals of physics* 303.1 (2003), pp. 2–30.
- [42] P.V. Buividovich and M.I. Polikarpov. “Entanglement entropy in gauge theories and the holographic principle for electric strings”. In: *Physics Letters B* 670.2 (2008), pp. 141–145. ISSN: 0370-2693. DOI: <https://doi.org/10.1016/j.physletb.2008.10.032>. URL: <https://www.sciencedirect.com/science/article/pii/S0370269308012987>.
- [43] William Donnelly. “Decomposition of entanglement entropy in lattice gauge theory”. In: *Phys. Rev. D* 85 (8 Apr. 2012), p. 085004. DOI: 10.1103/PhysRevD.85.085004. URL: <https://link.aps.org/doi/10.1103/PhysRevD.85.085004>.
- [44] Horacio Casini, Marina Huerta, and José Alejandro Rosabal. “Remarks on entanglement entropy for gauge fields”. In: *Phys. Rev. D* 89 (8 Apr. 2014), p. 085012. DOI: 10.1103/PhysRevD.89.085012. URL: <https://link.aps.org/doi/10.1103/PhysRevD.89.085012>.

- [45] Karel Van Acoleyen et al. “Entanglement of Distillation for Lattice Gauge Theories”. In: *Phys. Rev. Lett.* 117 (13 Sept. 2016), p. 131602. DOI: 10.1103/PhysRevLett.117.131602. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.117.131602>.
- [46] Ronak M. Soni and Sandip P. Trivedi. “Aspects of Entanglement Entropy for Gauge Theories”. In: *JHEP* 01 (2016), p. 136. DOI: 10.1007/JHEP01(2016)136. arXiv: 1510.07455 [hep-th]. URL: <https://arxiv.org/abs/1510.07455>.
- [47] Alexei Kitaev and John Preskill. “Topological Entanglement Entropy”. In: *Phys. Rev. Lett.* 96 (11 Mar. 2006), p. 110404. DOI: 10.1103/PhysRevLett.96.110404. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.96.110404>.
- [48] Michael Levin and Xiao-Gang Wen. “Detecting Topological Order in a Ground State Wave Function”. In: *Phys. Rev. Lett.* 96 (11 Mar. 2006), p. 110405. DOI: 10.1103/PhysRevLett.96.110405. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.96.110405>.
- [49] Scott Aaronson, Sean M. Carroll, and Lauren Ouellette. *Quantifying the Rise and Fall of Complexity in Closed Systems: The Coffee Automaton*. 2014. arXiv: 1405.6903 [quant-ph]. URL: <https://arxiv.org/abs/1405.6903>.
- [50] H.-G. Schöpf. “V. I. Arnold and A. Avez, Ergodic Problems of Classical Mechanics. (The Mathematical Physics Monograph Series) IX + 286 S. m. Fig. New York/Amsterdam 1968. W. A. Benjamin, Inc. Preis geb. \$ 14.75, brosch. \$ 6.95.” In: *ZAMM - Journal of Applied Mathematics and Mechanics / Zeitschrift für Angewandte Mathematik und Mechanik* 50.7-9 (1970), pp. 506–506. DOI: <https://doi.org/10.1002/zamm.19700500721>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/zamm.19700500721>. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/zamm.19700500721>.

- [51] Luca D'Alessio et al. "From quantum chaos and eigenstate thermalization to statistical mechanics and thermodynamics". In: *Advances in Physics* 65.3 (2016), pp. 239–362.
- [52] Josh M Deutsch. "Quantum statistical mechanics in a closed system". In: *Physical review a* 43.4 (1991), p. 2046.
- [53] Mark Srednicki. "Chaos and quantum thermalization". In: *Physical review e* 50.2 (1994), p. 888.
- [54] Hannes Bernien et al. "Probing many-body dynamics on a 51-atom quantum simulator". In: *Nature* 551.7682 (2017), pp. 579–584. DOI: <https://doi.org/10.1038/nature24622>.
- [55] Christopher J Turner et al. "Weak ergodicity breaking from quantum many-body scars". In: *Nature Physics* 14.7 (2018), pp. 745–749.
- [56] Maksym Serbyn, Dmitry A Abanin, and Zlatko Papić. "Quantum many-body scars and weak breaking of ergodicity". In: *Nature Physics* 17.6 (2021), pp. 675–685.
- [57] Sanjay Moudgalya, B Andrei Bernevig, and Nicolas Regnault. "Quantum many-body scars and Hilbert space fragmentation: a review of exact results". In: *Reports on Progress in Physics* 85.8 (2022), p. 086501.
- [58] Anushya Chandran et al. "Quantum many-body scars: A quasiparticle perspective". In: *Annual Review of Condensed Matter Physics* 14 (2023), pp. 443–469.
- [59] Seulgi Ok et al. "Topological many-body scar states in dimensions one, two, and three". In: *Physical Review Research* 1.3 (2019), p. 033144.
- [60] CJ Turner et al. "Quantum scarred eigenstates in a Rydberg atom chain: Entanglement, breakdown of thermalization, and stability to perturbations". In: *Physical Review B* 98.15 (2018), p. 155134.

- [61] Ian Mondragon-Shem, Maxim G Vavilov, and Ivar Martin. “Fate of quantum many-body scars in the presence of disorder”. In: *PRX Quantum* 2.3 (2021), p. 030349.
- [62] Cheng-Ju Lin, Vladimir Calvera, and Timothy H Hsieh. “Quantum many-body scar states in two-dimensional Rydberg atom arrays”. In: *Physical Review B* 101.22 (2020), p. 220304.
- [63] Bennet Windt and Hannes Pichler. “Squeezing quantum many-body scars”. In: *Physical Review Letters* 128.9 (2022), p. 090606.
- [64] Zhiyuan Yao et al. “Quantum many-body scars and quantum criticality”. In: *Physical Review B* 105.12 (2022), p. 125123.
- [65] I-Chi Chen et al. “Error-mitigated simulation of quantum many-body scars on quantum computers with pulse-level control”. In: *Physical Review Research* 4.4 (2022), p. 043027.
- [66] Eric J. Heller. “Bound-State Eigenfunctions of Classically Chaotic Hamiltonian Systems: Scars of Periodic Orbits”. In: *Phys. Rev. Lett.* 53 (16 Oct. 1984), pp. 1515–1518. DOI: 10.1103/PhysRevLett.53.1515. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.53.1515>.
- [67] Bart van Voorden, Jiří Minář, and Kareljan Schoutens. “Quantum many-body scars in transverse field Ising ladders and beyond”. In: *Physical Review B* 101.22 (2020), p. 220305.
- [68] Thomas Iadecola and Michael Schecter. “Quantum many-body scar states with emergent kinetic constraints and finite-entanglement revivals”. In: *Physical Review B* 101.2 (2020), p. 024306.
- [69] Kyungmin Lee et al. “Exact three-colored quantum scars from geometric frustration”. In: *Physical Review B* 101.24 (2020), p. 241111.

- [70] Naoyuki Shibata, Nobuyuki Yoshioka, and Hosho Katsura. “Onsager’s scars in disordered spin chains”. In: *Physical Review Letters* 124.18 (2020), p. 180604.
- [71] Sambuddha Chattopadhyay et al. “Quantum many-body scars from virtual entangled pairs”. In: *Physical Review B* 101.17 (2020), p. 174308.
- [72] Michael Schecter and Thomas Iadecola. “Weak Ergodicity Breaking and Quantum Many-Body Scars in Spin-1 X Y Magnets”. In: *Physical review letters* 123.14 (2019), p. 147201.
- [73] Soonwon Choi et al. “Emergent SU (2) dynamics and perfect quantum many-body scars”. In: *Physical review letters* 122.22 (2019), p. 220603.
- [74] Daniel K Mark, Cheng-Ju Lin, and Olexei I Motrunich. “Unified structure for exact towers of scar states in the Affleck-Kennedy-Lieb-Tasaki and other models”. In: *Physical Review B* 101.19 (2020), p. 195131.
- [75] Sanjay Moudgalya et al. “Large classes of quantum scarred Hamiltonians from matrix product states”. In: *Physical Review B* 102.8 (2020), p. 085120.
- [76] Sanjay Moudgalya, Nicolas Regnault, and B Andrei Bernevig. “Entanglement of exact excited states of Affleck-Kennedy-Lieb-Tasaki models: Exact results, many-body scars, and violation of the strong eigenstate thermalization hypothesis”. In: *Physical Review B* 98.23 (2018), p. 235156.
- [77] Sanjay Moudgalya et al. “Exact excited states of nonintegrable models”. In: *Physical Review B* 98.23 (2018), p. 235155.
- [78] NS Srivatsa et al. “Quantum many-body scars with chiral topological order in 2d and critical properties in 1d”. In: *arXiv preprint arXiv:2009.03923* ().
- [79] Naoto Shiraishi. “Connection between quantum-many-body scars and the Affleck–Kennedy–Lieb–Tasaki model from the viewpoint of embedded Hamiltonians”. In: *Journal of Statistical Mechanics: Theory and Experiment* 2019.8 (2019), p. 083103.

- [80] Julia Wildeboer et al. “Topological quantum many-body scars in quantum dimer models on the kagome lattice”. In: *Physical Review B* 104.12 (2021), p. L121103.
- [81] Christopher M Langlett et al. “Rainbow scars: From area to volume law”. In: *Physical Review B* 105.6 (2022), p. L060301.
- [82] Jean-Yves Desaulès et al. “Prominent quantum many-body scars in a truncated Schwinger model”. In: *Physical Review B* 107.20 (2023), p. 205112.
- [83] Giuseppe Calajó et al. “Quantum many-body scarring in a non-Abelian lattice gauge theory”. In: *Phys. Rev. Res.* 7 (1 Mar. 2025), p. 013322. DOI: 10.1103/PhysRevResearch.7.013322. URL: <https://link.aps.org/doi/10.1103/PhysRevResearch.7.013322>.
- [84] Thea Budde, Marina Krstic Marinkovic, and Joao C Pinto Barros. “Quantum many-body scars for arbitrary integer spin in 2+ 1 D Abelian gauge theories”. In: *Physical Review D* 110.9 (2024), p. 094506.
- [85] Jesse Osborne, Ian P McCulloch, and Jad C Halimeh. “Quantum Many-Body Scarring in 2+1 D Gauge Theories with Dynamical Matter”. In: *arXiv preprint arXiv:2403.08858* (2024).
- [86] Giuseppe Calajò et al. “Quantum many-body scarring in a non-Abelian lattice gauge theory”. In: *arXiv preprint arXiv:2405.13112* (2024).
- [87] Federica M. Surace et al. “Lattice Gauge Theories and String Dynamics in Rydberg Atom Quantum Simulators”. In: *Phys. Rev. X* 10 (2 May 2020), p. 021041. DOI: 10.1103/PhysRevX.10.021041. URL: <https://link.aps.org/doi/10.1103/PhysRevX.10.021041>.
- [88] Kiryl Pakrouski et al. “Many-body scars as a group invariant sector of Hilbert space”. In: *Physical review letters* 125.23 (2020), p. 230602.

- [89] Sanjay Moudgalya and Olexei I. Motrunich. “Hilbert Space Fragmentation and Commutant Algebras”. In: *Phys. Rev. X* 12 (1 Mar. 2022), p. 011050. DOI: 10.1103/PhysRevX.12.011050. URL: <https://link.aps.org/doi/10.1103/PhysRevX.12.011050>.
- [90] Sanjay Moudgalya and Olexei I. Motrunich. “From symmetries to commutant algebras in standard Hamiltonians”. In: *Annals of Physics* 455 (2023), p. 169384. ISSN: 0003-4916. DOI: <https://doi.org/10.1016/j.aop.2023.169384>. URL: <https://www.sciencedirect.com/science/article/pii/S0003491623001707>.
- [91] Zimo Sun et al. “Majorana scars as group singlets”. In: *Physical Review Research* 5.4 (2023), p. 043208.
- [92] Sanjay Moudgalya, Nicolas Regnault, and B. Andrei Bernevig. “ η -pairing in Hubbard models: From spectrum generating algebras to quantum many-body scars”. In: *Phys. Rev. B* 102 (8 Aug. 2020), p. 085140. DOI: 10.1103/PhysRevB.102.085140. URL: <https://link.aps.org/doi/10.1103/PhysRevB.102.085140>.
- [93] Daniel Gottesman. “The Heisenberg representation of quantum computers”. In: *arXiv preprint quant-ph/9807006* (1998).
- [94] Daniel M. Greenberger, Michael A. Horne, and Anton Zeilinger. “Going Beyond Bell’s Theorem”. In: *Bell’s Theorem, Quantum Theory, and Conceptions of the Universe*. Ed. by Menas Kafatos. Vol. 37. Fundamental Theories of Physics. Dordrecht: Kluwer Academic, 1989, pp. 69–72.
- [95] Scott Aaronson and Daniel Gottesman. “Improved simulation of stabilizer circuits”. In: *Physical Review A—Atomic, Molecular, and Optical Physics* 70.5 (2004), p. 052328.

- [96] Sergey Bravyi and Alexei Kitaev. “Universal quantum computation with ideal Clifford gates and noisy ancillas”. In: *Physical Review A—Atomic, Molecular, and Optical Physics* 71.2 (2005), p. 022316.
- [97] Earl T Campbell, Barbara M Terhal, and Christophe Vuillot. “Roads towards fault-tolerant universal quantum computation”. In: *Nature* 549.7671 (2017), pp. 172–179.
- [98] Lorenzo Leone, Salvatore FE Oliviero, and Alioscia Hama. “Stabilizer rényi entropy”. In: *Physical Review Letters* 128.5 (2022), p. 050402.
- [99] Xhek Turkeshi, Anatoly Dymarsky, and Piotr Sierant. “Pauli spectrum and magic of typical quantum many-body states”. In: *arXiv preprint arXiv:2312.11631* (2023).
- [100] Zi-Wen Liu and Andreas Winter. “Many-body quantum magic”. In: *PRX Quantum* 3.2 (2022), p. 020333.
- [101] Ryan Smith, Zlatko Papić, and Andrew Hallam. “Non-stabilizerness in kinetically-constrained Rydberg atom arrays”. In: *arXiv preprint arXiv:2406.14348* (2024).
- [102] Sabhyata Gupta et al. “Exact stabilizer scars in two-dimensional U (1) lattice gauge theory”. In: *Physical Review D* 113.9 (2026), p. 094509.
- [103] Shane Dooley. “Parent Hamiltonians for stabilizer quantum many-body scars”. In: *arXiv preprint arXiv:2601.10805* (2026).
- [104] Jens Eisert et al. “Quantum certification and benchmarking”. In: *Nature Reviews Physics* 2.7 (2020), pp. 382–390. DOI: 10.1038/s42254-020-0186-4. URL: <https://doi.org/10.1038/s42254-020-0186-4>.
- [105] Daniel F. V. James et al. “Measurement of qubits”. In: *Phys. Rev. A* 64 (5 Oct. 2001), p. 052312. DOI: 10.1103/PhysRevA.64.052312. URL: <https://link.aps.org/doi/10.1103/PhysRevA.64.052312>.

- [106] Jeongwan Haah et al. “Sample-Optimal Tomography of Quantum States”. In: *IEEE Transactions on Information Theory* 63.9 (2017), pp. 5628–5641. DOI: 10.1109/TIT.2017.2719044.
- [107] Joseph Emerson, Robert Alicki, and Karol Życzkowski. “Scalable noise estimation with random unitary operators”. In: *Journal of Optics B: Quantum and Semiclassical Optics* 7.10 (Sept. 2005), S347. DOI: 10.1088/1464-4266/7/10/021. URL: <https://doi.org/10.1088/1464-4266/7/10/021>.
- [108] E. Knill et al. “Randomized benchmarking of quantum gates”. In: *Phys. Rev. A* 77 (1 Jan. 2008), p. 012307. DOI: 10.1103/PhysRevA.77.012307. URL: <https://link.aps.org/doi/10.1103/PhysRevA.77.012307>.
- [109] Easwar Magesan, J. M. Gambetta, and Joseph Emerson. “Scalable and Robust Randomized Benchmarking of Quantum Processes”. In: *Phys. Rev. Lett.* 106 (18 May 2011), p. 180504. DOI: 10.1103/PhysRevLett.106.180504. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.106.180504>.
- [110] Steven T Flammia and Yi-Kai Liu. “Direct fidelity estimation from few Pauli measurements”. In: *Physical review letters* 106.23 (2011), p. 230501.
- [111] Marcus P da Silva, Olivier Landon-Cardinal, and David Poulin. “Practical characterization of quantum devices without tomography”. In: *Physical Review Letters* 107.21 (2011), p. 210404.
- [112] Lorenzo Leone, Salvatore F. E. Oliviero, and Alioscia Hamma. “Nonstabilizerness determining the hardness of direct fidelity estimation”. In: *Phys. Rev. A* 107 (2 Feb. 2023), p. 022429. DOI: 10.1103/PhysRevA.107.022429. URL: <https://link.aps.org/doi/10.1103/PhysRevA.107.022429>.
- [113] Jeremy Hartse and Alessandro Roggero. “Faster spectral density calculation using energy moments”. In: *The European Physical Journal A* 59.41 (2023). DOI: 10.1140/epja/s10050-023-00952-6. URL: <https://doi.org>.

- [114] V D Efros et al. “The Lorentz integral transform (LIT) method and its applications to perturbation-induced reactions”. In: *Journal of Physics G: Nuclear and Particle Physics* 34.12 (Oct. 2007), R459–R528. ISSN: 1361-6471. DOI: 10.1088/0954-3899/34/12/r02. URL: <http://dx.doi.org/10.1088/0954-3899/34/12/R02>.
- [115] A. Roggero. “Spectral-density estimation with the Gaussian integral transform”. In: *Phys. Rev. A* 102 (2 Aug. 2020), p. 022409. DOI: 10.1103/PhysRevA.102.022409. URL: <https://link.aps.org/doi/10.1103/PhysRevA.102.022409>.
- [116] Joanna E. Sobczyk and Alessandro Roggero. “Spectral density reconstruction with Chebyshev polynomials”. In: *Phys. Rev. E* 105 (5 May 2022), p. 055310. DOI: 10.1103/PhysRevE.105.055310. URL: <https://link.aps.org/doi/10.1103/PhysRevE.105.055310>.
- [117] W. Glöckle and M. Schwamb. “On the Ill-Posed Character of the Lorentz Integral Transform”. In: *Few-Body Systems* 46.1 (Mar. 2009), pp. 55–62. ISSN: 1432-5411. DOI: 10.1007/s00601-009-0050-7. URL: <http://dx.doi.org/10.1007/s00601-009-0050-7>.
- [118] N. Barnea et al. “The Lorentz Integral Transform and its Inversion”. In: *Few-Body Systems* 47.4 (Jan. 2010), pp. 201–206. ISSN: 1432-5411. DOI: 10.1007/s00601-009-0081-0. URL: <http://dx.doi.org/10.1007/s00601-009-0081-0>.
- [119] R. N. Silver, D. S. Sivia, and J. E. Gubernatis. “Maximum-entropy method for analytic continuation of quantum Monte Carlo data”. In: *Phys. Rev. B* 41 (4 Feb. 1990), pp. 2380–2389. DOI: 10.1103/PhysRevB.41.2380. URL: <https://link.aps.org/doi/10.1103/PhysRevB.41.2380>.

- [120] E. Vitali et al. “Ab initio low-energy dynamics of superfluid and solid ^4He ”. In: *Phys. Rev. B* 82 (17 Nov. 2010), p. 174510. DOI: 10.1103/PhysRevB.82.174510. URL: <https://link.aps.org/doi/10.1103/PhysRevB.82.174510>.
- [121] Yannis Burnier and Alexander Rothkopf. “Bayesian Approach to Spectral Function Reconstruction for Euclidean Quantum Field Theories”. In: *Phys. Rev. Lett.* 111 (18 Oct. 2013), p. 182003. DOI: 10.1103/PhysRevLett.111.182003. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.111.182003>.
- [122] Lukas Kades et al. “Spectral reconstruction with deep neural networks”. In: *Phys. Rev. D* 102 (9 Nov. 2020), p. 096001. DOI: 10.1103/PhysRevD.102.096001. URL: <https://link.aps.org/doi/10.1103/PhysRevD.102.096001>.
- [123] Krishnan Raghavan et al. “Machine-learning-based inversion of nuclear responses”. In: *Phys. Rev. C* 103 (3 Mar. 2021), p. 035502. DOI: 10.1103/PhysRevC.103.035502. URL: <https://link.aps.org/doi/10.1103/PhysRevC.103.035502>.
- [124] J. E. Sobczyk et al. “Spectral function for ^4He using the Chebyshev expansion in coupled-cluster theory”. In: *Phys. Rev. C* 106 (3 Sept. 2022), p. 034310. DOI: 10.1103/PhysRevC.106.034310. URL: <https://link.aps.org/doi/10.1103/PhysRevC.106.034310>.
- [125] Andrew M. Childs, Robin Kothari, and Rolando D. Somma. “Quantum Algorithm for Systems of Linear Equations with Exponentially Improved Dependence on Precision”. In: *SIAM Journal on Computing* 46.6 (2017), pp. 1920–1950. DOI: 10.1137/16M1087072. URL: <https://doi.org/10.1137/16M1087072>.
- [126] Guang Hao Low and Isaac L. Chuang. “Hamiltonian Simulation by Qubitization”. In: *Quantum* 3 (July 2019), p. 163. ISSN: 2521-327X. DOI: 10.22331/q-2019-07-12-163. URL: <https://doi.org/10.22331/q-2019-07-12-163>.

- [127] András Gilyén et al. “Quantum singular value transformation and beyond: exponential improvements for quantum matrix arithmetics”. In: *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*. 2019, pp. 193–204.
- [128] Sathyawageeswar Subramanian, Stephen Brierley, and Richard Jozsa. “Implementing smooth functions of a Hermitian matrix on a quantum computer”. In: *Journal of Physics Communications* 3.6 (June 2019), p. 065002. DOI: 10.1088/2399-6528/ab25a2. URL: <https://doi.org/10.1088/2399-6528/ab25a2>.
- [129] G Orlandini and M Traini. “Sum rules for electron-nucleus scattering”. In: *Reports on Progress in Physics* 54.2 (Feb. 1991), pp. 257–338. DOI: 10.1088/0034-4885/54/2/002. URL: <https://doi.org/10.1088/0034-4885/54/2/002>.
- [130] R Rosenfelder. “Quasielastic electron scattering from nuclei”. In: *Annals of Physics* 128.1 (1980), pp. 188–240. ISSN: 0003-4916. DOI: [https://doi.org/10.1016/0003-4916\(80\)90059-7](https://doi.org/10.1016/0003-4916(80)90059-7). URL: <https://www.sciencedirect.com/science/article/pii/0003491680900597>.
- [131] J. Carlson and R. Schiavilla. “Structure and dynamics of few-nucleon systems”. In: *Rev. Mod. Phys.* 70 (3 July 1998), pp. 743–841. DOI: 10.1103/RevModPhys.70.743. URL: <https://link.aps.org/doi/10.1103/RevModPhys.70.743>.
- [132] J. Carlson and R. Schiavilla. “Euclidean proton response in light nuclei”. In: *Phys. Rev. Lett.* 68 (25 June 1992), pp. 3682–3685. DOI: 10.1103/PhysRevLett.68.3682. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.68.3682>.
- [133] J. Carlson et al. “Quantum Monte Carlo methods for nuclear physics”. In: *Rev. Mod. Phys.* 87 (3 Sept. 2015), pp. 1067–1118. DOI: 10.1103/RevModPhys.87.1067. URL: <https://link.aps.org/doi/10.1103/RevModPhys.87.1067>.

- [134] Thomas E O’Brien, Brian Tarasinski, and Barbara M Terhal. “Quantum phase estimation of multiple eigenvalues for small-scale (noisy) experiments”. In: *New Journal of Physics* 21.2 (Feb. 2019), p. 023022. DOI: 10.1088/1367-2630/aafb8e. URL: <https://doi.org/10.1088/1367-2630/aafb8e>.
- [135] Rolando D Somma. “Quantum eigenvalue estimation via time series analysis”. In: *New Journal of Physics* 21.12 (Dec. 2019), p. 123025. DOI: 10.1088/1367-2630/ab5c60. URL: <https://doi.org/10.1088/1367-2630/ab5c60>.
- [136] Sirui Lu, Mari Carmen Bañuls, and J. Ignacio Cirac. “Algorithms for Quantum Simulation at Finite Energies”. In: *PRX Quantum* 2 (2 May 2021), p. 020321. DOI: 10.1103/PRXQuantum.2.020321. URL: <https://link.aps.org/doi/10.1103/PRXQuantum.2.020321>.
- [137] Edgar Andres Ruiz Guzman and Denis Lacroix. *Calculation of generating function in many-body systems with quantum computers: technical challenges and use in hybrid quantum-classical methods*. 2021. DOI: 10.48550/ARXIV.2104.08181. URL: <https://arxiv.org/abs/2104.08181>.
- [138] Cristian L. Cortes and Stephen K. Gray. “Quantum Krylov subspace algorithms for ground- and excited-state energy estimation”. In: *Phys. Rev. A* 105 (2 Feb. 2022), p. 022417. DOI: 10.1103/PhysRevA.105.022417. URL: <https://link.aps.org/doi/10.1103/PhysRevA.105.022417>.
- [139] Alexander Weiße et al. “The kernel polynomial method”. In: *Reviews of Modern Physics* 78.1 (Mar. 2006), pp. 275–306. ISSN: 1539-0756. DOI: 10.1103/revmodphys.78.275. URL: <http://dx.doi.org/10.1103/RevModPhys.78.275>.
- [140] A. Baroni et al. “Nuclear two point correlation functions on a quantum computer”. In: *Phys. Rev. D* 105 (7 Apr. 2022), p. 074503. DOI: 10.1103/PhysRevD.105.074503. URL: <https://link.aps.org/doi/10.1103/PhysRevD.105.074503>.

- [141] J. Ahrens et al. “Total nuclear photon absorption cross sections for some light elements”. In: *Nuclear Physics A* 251.3 (1975), pp. 479–492. ISSN: 0375-9474. DOI: 10.1016/0375-9474(75)90543-6. URL: <https://www.sciencedirect.com/science/article/pii/0375947475905436>.
- [142] S. Bacca et al. “Giant and pigmy dipole resonances in ^4He , $^{16,22}\text{O}$, and ^{40}Ca from chiral nucleon-nucleon interactions”. In: *Phys. Rev. C* 90 (6 Dec. 2014), p. 064619. DOI: 10.1103/PhysRevC.90.064619. URL: <https://link.aps.org/doi/10.1103/PhysRevC.90.064619>.
- [143] C. F. Williamson et al. “Quasielastic electron scattering from ^{40}Ca ”. In: *Phys. Rev. C* 56 (6 Dec. 1997), pp. 3152–3172. DOI: 10.1103/PhysRevC.56.3152. URL: <https://link.aps.org/doi/10.1103/PhysRevC.56.3152>.
- [144] J. E. Sobczyk et al. “Ab Initio Computation of the Longitudinal Response Function in ^{40}Ca ”. In: *Phys. Rev. Lett.* 127 (7 Aug. 2021), p. 072501. DOI: 10.1103/PhysRevLett.127.072501. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.127.072501>.
- [145] Shouzheng Gu, Rolando D. Somma, and Burak Şahinoğlu. “Fast-forwarding quantum evolution”. In: *Quantum* 5 (Nov. 2021), p. 577. ISSN: 2521-327X. DOI: 10.22331/q-2021-11-15-577. URL: <https://doi.org/10.22331/q-2021-11-15-577>.
- [146] Andrew M. Childs and Nathan Wiebe. “Hamiltonian simulation using linear combinations of unitary operations”. In: *Quantum Information and Computation* 12 (1112 2012), pp. 0901–0924. DOI: 10.26421/QIC12.11-12. URL: <http://www.rintonpress.com/xxqic12/qic-12-1112/0901-0924.pdf>.
- [147] Yimin Ge, Jordi Tura, and J. Ignacio Cirac. “Faster ground state preparation and high-precision ground energy estimation with fewer qubits”. In: *Journal of Mathematical Physics* 60.2 (2019), p. 022202. DOI: 10.1063/1.5027484. URL: <https://doi.org/10.1063/1.5027484>.

- [148] Thomas E. O’Brien et al. “Error Mitigation via Verified Phase Estimation”. In: *PRX Quantum* 2 (2 May 2021), p. 020317. DOI: 10.1103/PRXQuantum.2.020317. URL: <https://link.aps.org/doi/10.1103/PRXQuantum.2.020317>.
- [149] Jeremy Hartse, Lukasz Fidkowski, and Niklas Mueller. “Stabilizer scars”. In: *Physical Review Letters* 135.6 (2025), p. 060402.
- [150] Christian Gross and Immanuel Bloch. “Quantum simulations with ultracold atoms in optical lattices”. In: *Science* 357.6355 (2017), pp. 995–1001.
- [151] Jens Eisert, Mathis Friesdorf, and Christian Gogolin. “Quantum many-body systems out of equilibrium”. In: *Nature Physics* 11.2 (2015), pp. 124–130.
- [152] Masahito Ueda. “Quantum equilibration, thermalization and prethermalization in ultracold atoms”. In: *Nature Reviews Physics* 2.12 (2020), pp. 669–681.
- [153] Jürgen Berges et al. “QCD thermalization: Ab initio approaches and interdisciplinary connections”. In: *Reviews of Modern Physics* 93.3 (2021), p. 035003.
- [154] Thomas Iadecola and Marko Žnidarič. “Exact localized and ballistic eigenstates in disordered chaotic spin ladders and the Fermi-Hubbard model”. In: *Physical Review Letters* 123.3 (2019), p. 036403.
- [155] Andrea Pizzi et al. “Quantum scars in many-body systems”. In: *arXiv preprint arXiv:2408.10301* (2024).
- [156] Natalie Klco, Alessandro Roggero, and Martin J. Savage. “Standard model physics and the digital quantum revolution: thoughts about the interface”. In: *Rept. Prog. Phys.* 85.6 (2022), p. 064301. DOI: 10.1088/1361-6633/ac58a4. arXiv: 2107.04769 [quant-ph].
- [157] Christian W Bauer et al. “Quantum simulation of fundamental particles and forces”. In: *Nature Reviews Physics* 5.7 (2023), pp. 420–432.

- [158] Niklas Mueller, Torsten V Zache, and Robert Ott. “Thermalization of gauge theories from their entanglement spectrum”. In: *Physical Review Letters* 129.1 (2022), p. 011601.
- [159] Zhao-Yu Zhou et al. “Thermalization dynamics of a gauge theory on a quantum simulator”. In: *Science* 377.6603 (2022), pp. 311–314.
- [160] Lukas Ebner et al. “Eigenstate thermalization in $(2+1)$ -dimensional $SU(2)$ lattice gauge theory”. In: *Physical Review D* 109.1 (2024), p. 014504.
- [161] Niklas Mueller et al. “Quantum Computing Universal Thermalization Dynamics in a $(2+1)$ D Lattice Gauge Theory”. In: *arXiv preprint arXiv:2408.00069* (2024).
- [162] Eduardo Fradkin. *Field theories of condensed matter physics*. Cambridge University Press, 2013.
- [163] Hagen Kleinert. *Gauge Fields in Condensed Matter: Vol. 1: Superflow and Vortex Lines (Disorder Fields, Phase Transitions) Vol. 2: Stresses and Defects (Differential Geometry, Crystal Melting)*. World Scientific, 1989.
- [164] Sankar Das Sarma, Michael Freedman, and Chetan Nayak. “Topological quantum computation”. In: *Physics Today* 59.7 (2006), pp. 32–38. URL: <https://doi.org/10.1063/1.2337825>.
- [165] Chetan Nayak et al. “Non-Abelian anyons and topological quantum computation”. In: *Rev. Mod. Phys.* 80 (3 Sept. 2008), pp. 1083–1159. DOI: 10.1103/RevModPhys.80.1083. URL: <https://link.aps.org/doi/10.1103/RevModPhys.80.1083>.
- [166] Ville Lahtinen and Jiannis K. Pachos. “A Short Introduction to Topological Quantum Computation”. In: *SciPost Phys.* 3 (2017), p. 021. DOI: 10.21468/SciPostPhys.3.3.021. URL: <https://scipost.org/10.21468/SciPostPhys.3.3.021>.

- [167] Franz J Wegner. “Duality in generalized Ising models and phase transitions without local order parameters”. In: *Journal of Mathematical Physics* 12.10 (1971), pp. 2259–2272.
- [168] David Horn, Marvin Weinstein, and Shimon Yankielowicz. “Hamiltonian approach to $Z(N)$ lattice gauge theories”. In: *Physical Review D* 19.12 (1979), p. 3715.
- [169] Subir Sachdev. “Topological order, emergent gauge fields, and Fermi surface reconstruction”. In: *Reports on Progress in Physics* 82.1 (2018), p. 014001.
- [170] Sanjay Moudgalya and Olexei I. Motrunich. “Hilbert Space Fragmentation and Commutant Algebras”. In: *Physical Review X* 12.1 (Mar. 2022). ISSN: 2160-3308. DOI: 10.1103/physrevx.12.011050.
- [171] Sanjay Moudgalya and Olexei I. Motrunich. *Exhaustive Characterization of Quantum Many-Body Scars using Commutant Algebras*. 2023. arXiv: 2209.03377 [cond-mat.str-el]. URL: <https://arxiv.org/abs/2209.03377>.
- [172] Sanjay Moudgalya and Olexei I. Motrunich. “From symmetries to commutant algebras in standard Hamiltonians”. In: *Annals of Physics* 455 (Aug. 2023), p. 169384. ISSN: 0003-4916. DOI: 10.1016/j.aop.2023.169384. URL: <http://dx.doi.org/10.1016/j.aop.2023.169384>.
- [173] Bart van Voorden, Jiří Minář, and Kareljan Schoutens. “Quantum many-body scars in transverse field Ising ladders and beyond”. In: *Phys. Rev. B* 101 (22 June 2020), p. 220305. DOI: 10.1103/PhysRevB.101.220305. URL: <https://link.aps.org/doi/10.1103/PhysRevB.101.220305>.
- [174] Rolando Somma et al. “Simulating physical phenomena by quantum networks”. In: *Physical Review A* 65.4 (2002), p. 042323.

- [175] Michael Knap et al. “Probing real-space and time-resolved correlation functions with many-body Ramsey interferometry”. In: *Physical review letters* 111.14 (2013), p. 147205.
- [176] Niklas Mueller et al. “Quantum computation of dynamical quantum phase transitions and entanglement tomography in a lattice gauge theory”. In: *PRX Quantum* 4.3 (2023), p. 030323.
- [177] Andreas Elben et al. “The randomized measurement toolbox”. In: *Nat. Rev. Phys.* 5.1 (2023), pp. 9–24. URL: <https://doi.org/10.1038/s42254-022-00535-2>.
- [178] Jacob Bringewatt, Jonathan Kunjummen, and Niklas Mueller. “Randomized measurement protocols for lattice gauge theories”. In: *Quantum* 8 (2024), p. 1300.
- [179] Christian Kokail et al. “Entanglement Hamiltonian tomography in quantum simulation”. In: *Nat. Phys.* 17.8 (2021), pp. 936–942. URL: <https://www.nature.com/articles/s41567-021-01260-w>.
- [180] Christian Kokail et al. “Quantum variational learning of the entanglement hamiltonian”. In: *Phys. Rev. Lett.* 127.17 (2021), p. 170501. URL: <https://doi.org/10.1103/PhysRevLett.127.170501>.
- [181] Poetri Sonya Tarabunga et al. “Many-body magic via pauli-markov chains—from criticality to gauge theories”. In: *PRX Quantum* 4.4 (2023), p. 040317.
- [182] Davide Rattacaso et al. “Stabilizer entropy dynamics after a quantum quench”. In: *Physical Review A* 108.4 (2023), p. 042407.
- [183] AG Catalano et al. “Magic phase transition and non-local complexity in generalized W State”. In: *arXiv preprint arXiv:2406.19457* (2024).
- [184] Caroline EP Robin and Martin J Savage. “The Magic in Nuclear and Hypernuclear Forces”. In: *arXiv preprint arXiv:2405.10268* (2024).

- [185] Florian Brökemeier et al. “Quantum Magic and Multi-Partite Entanglement in the Structure of Nuclei”. In: *arXiv preprint arXiv:2409.12064* (2024).
- [186] Ivan Chernyshev, Caroline EP Robin, and Martin J Savage. “Quantum Magic and Computational Complexity in the Neutrino Sector”. In: *arXiv preprint arXiv:2411.04203* (2024).
- [187] J Hartse, F Lukasz, and N Mueller. *Stabilizer Scars Data*. 2025. DOI: 10.5281/zenodo.16293590.
- [188] Saptarshi Biswas, Debasish Banerjee, and Arnab Sen. “Scars from protected zero modes and beyond in $U(1)$ quantum link and quantum dimer models”. In: *SciPost Physics* 12.5 (2022), p. 148.
- [189] Juan M Pino et al. “Demonstration of the trapped-ion quantum CCD computer architecture”. In: *Nature* 592.7853 (2021), pp. 209–213.
- [190] Kenneth Rudinger et al. “Characterizing midcircuit measurements on a superconducting qubit using gate set tomography”. In: *Physical Review Applied* 17.1 (2022), p. 014014.
- [191] Ludmila Botelho et al. “Error mitigation for variational quantum algorithms through mid-circuit measurements”. In: *Physical Review A* 105.2 (2022), p. 022441.
- [192] Hans J Briegel et al. “Measurement-based quantum computation”. In: *Nature Physics* 5.1 (2009), pp. 19–26.
- [193] Miguel Aguado et al. “Creation, manipulation, and detection of Abelian and non-Abelian anyons in optical lattices”. In: *Physical review letters* 101.26 (2008), p. 260501.
- [194] Nathanan Tantivasadakarn et al. “Long-range entanglement from measuring symmetry-protected topological phases”. In: *Physical Review X* 14.2 (2024), p. 021040.

- [195] Elisa Bäumer and Stefan Woerner. “Measurement-Based Long-Range Entangling Gates in Constant Depth”. In: *arXiv preprint arXiv:2408.03064* (2024).
- [196] Scholarpedia http://www.scholarpedia.org/article/Loschmidt_echo (Oct 31, 2024).
- [197] Armando Angrisani et al. “Classically estimating observables of noiseless quantum circuits”. In: *arXiv preprint arXiv:2409.01706* (2024).
- [198] Jeremy Hartse et al. *Benchmarking quantum simulation at scale*. 2026. arXiv: 2607.14212 [quant-ph].
- [199] Scott Aaronson and Lijie Chen. “Complexity-theoretic foundations of quantum supremacy experiments”. In: *arXiv preprint arXiv:1612.05903* (2016).
- [200] Ivan H Deutsch. “Harnessing the power of the second quantum revolution”. In: *PRX Quantum* 1.2 (2020), p. 020101.
- [201] Peter W Shor. “Algorithms for quantum computation: discrete logarithms and factoring”. In: *Proceedings 35th annual symposium on foundations of computer science*. Ieee. 1994, pp. 124–134.
- [202] Amira Abbas et al. “Challenges and opportunities in quantum optimization”. In: *Nature Reviews Physics* (2024), pp. 1–18.
- [203] Dan Shepherd and Michael J Bremner. “Temporally unstructured quantum computation”. In: *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 465.2105 (2009), pp. 1413–1439.
- [204] Scott Aaronson and Alex Arkhipov. “The computational complexity of linear optics”. In: *Proceedings of the forty-third annual ACM symposium on Theory of computing*. 2011, pp. 333–342.
- [205] Michael J Bremner, Ashley Montanaro, and Dan J Shepherd. “Achieving quantum supremacy with sparse and noisy commuting quantum computations”. In: *Quantum* 1 (2017), p. 8.

- [206] Sergio Boixo et al. “Characterizing quantum supremacy in near-term devices”. In: *Nature Physics* 14.6 (2018), pp. 595–600.
- [207] J Ignacio Cirac and Peter Zoller. “Goals and opportunities in quantum simulation”. In: *Nature physics* 8.4 (2012), pp. 264–266.
- [208] D Beck et al. “Nuclear physics and quantum information science”. In: *Report from the NSAC QIS Subcommittee* (2019).
- [209] Yudong Cao et al. “Quantum chemistry in the age of quantum computing”. In: *Chemical reviews* 119.19 (2019), pp. 10856–10915.
- [210] Alexei Yu Kitaev, Alexander Shen, and Mikhail N Vyalyi. *Classical and quantum computation*. 47. American Mathematical Soc., 2002.
- [211] Joseph F Fitzsimons, Michal Hajdušek, and Tomoyuki Morimae. “Post hoc verification of quantum computation”. In: *Physical review letters* 120.4 (2018), p. 040501.
- [212] Andreas Elben et al. “Cross-platform verification of intermediate scale quantum devices”. In: *Physical review letters* 124.1 (2020), p. 010504.
- [213] Daniel Mills et al. “Application-motivated, holistic benchmarking of a full quantum computing stack”. In: *Quantum* 5 (2021), p. 415.
- [214] Martin Kliesch and Ingo Roth. “Theory of quantum system certification”. In: *PRX quantum* 2.1 (2021), p. 010201.
- [215] Scott Aaronson and Yuxuan Zhang. “On verifiable quantum advantage with peaked circuit sampling”. In: *arXiv preprint arXiv:2404.14493* (2024).
- [216] Timothy Proctor et al. “Benchmarking quantum computers”. In: *Nature Reviews Physics* 7.2 (2025), pp. 105–118.
- [217] Peter F Kolb and Ulrich Heinz. “Hydrodynamic description of ultrarelativistic heavy-ion collisions”. In: *Quark-Gluon Plasma 3*. World Scientific, 2004, pp. 634–714.

- [218] J Ignacio Cirac and Frank Verstraete. “Renormalization and tensor product states in spin chains and lattices”. In: *Journal of physics a: mathematical and theoretical* 42.50 (2009), p. 504004.
- [219] Matt Probert. *Electronic Structure: Basic Theory and Practical Methods*, by Richard M. Martin: Scope: graduate level textbook. Level: theoretical materials scientists/condensed matter physicists/computational chemists. 2011.
- [220] Dorit Aharonov. “Quantum to classical phase transition in noisy quantum computers”. In: *Physical Review A* 62.6 (2000), p. 062311.
- [221] Aram W Harrow and Michael A Nielsen. “Robustness of quantum gates in the presence of noise”. In: *Physical Review A* 68.1 (2003), p. 012308.
- [222] Stuart Flannigan et al. “Propagation of errors and quantitative quantum simulation with quantum advantage”. In: *Quantum Science & Technology* 7.4 (2022), p. 045025.
- [223] Thomas Schuster et al. “A polynomial-time classical algorithm for noisy quantum circuits”. In: *Physical Review X* 15.4 (2025), p. 041018.
- [224] Armando Angrisani et al. “Simulating quantum circuits with arbitrary local noise using pauli propagation”. In: *arXiv preprint arXiv:2501.13101* (2025).
- [225] Shravan Shravan, Mohsin Raza, and Ariel Shlosberg. “Efficient simulation of noisy IQP circuits with amplitude-damping noise”. In: *arXiv preprint arXiv:2604.05036* (2026).
- [226] Rahul Trivedi, Adrian Franco Rubio, and J Ignacio Cirac. “Quantum advantage and stability to errors in analogue quantum simulators”. In: *Nature Communications* 15.1 (2024), p. 6507.
- [227] Zhenning Liu et al. “Efficiently verifiable quantum advantage on near-term analog quantum simulators”. In: *PRX Quantum* 6.1 (2025), p. 010341.

- [228] Matthew DeCross et al. “Computational power of random quantum circuits in arbitrary geometries”. In: *Physical Review X* 15.2 (2025), p. 021052.
- [229] Dongxin Gao et al. “Establishing a new benchmark in quantum computational advantage with 105-qubit Zuchongzhi 3.0 processor”. In: *Physical Review Letters* 134.9 (2025), p. 090601.
- [230] Alexis Morvan et al. “Phase transitions in random circuit sampling”. In: *Nature* 634.8033 (2024), pp. 328–333.
- [231] Yulin Wu et al. “Strong quantum computational advantage using a superconducting quantum processor”. In: *Physical review letters* 127.18 (2021), p. 180501.
- [232] Timothy J Proctor et al. “Direct randomized benchmarking for multiqubit devices”. In: *Physical review letters* 123.3 (2019), p. 030503.
- [233] Andrew W Cross et al. “Validating quantum computers using randomized model circuits”. In: *Physical Review A* 100.3 (2019), p. 032328.
- [234] Erik Nielsen et al. “Gate set tomography”. In: *Quantum* 5 (2021), p. 557.
- [235] Jonas Helsen et al. “General framework for randomized benchmarking”. In: *PRX quantum* 3.2 (2022), p. 020357.
- [236] Robin Blume-Kohout and Kevin C. Young. “A volumetric framework for quantum computer benchmarks”. In: *Quantum* 4 (Nov. 2020), p. 362. ISSN: 2521-327X. DOI: 10.22331/q-2020-11-15-362. URL: <https://doi.org/10.22331/q-2020-11-15-362>.
- [237] Jad C Halimeh et al. “Quantum simulation of out-of-equilibrium dynamics in gauge theories”. In: *arXiv preprint arXiv:2509.03586* (2025).

- [238] Daniel Gottesman. “An introduction to quantum error correction and fault-tolerant quantum computation”. In: *Quantum information science and its contributions to mathematics, Proceedings of Symposia in Applied Mathematics*. Vol. 68. 2010, pp. 13–58.
- [239] Akihiro Hokkyo. “Exact Thermal Stabilizer Eigenstates at Infinite Temperature”. In: *arXiv e-prints* (2026), arXiv–2601.
- [240] Tobias Haug and Lorenzo Piroli. “Stabilizer entropies and nonstabilizerness monotones”. In: *Quantum* 7 (2023), p. 1092.
- [241] Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. *Concentration Inequalities: A Nonasymptotic Theory of Independence*. Oxford University Press, 2013.
- [242] Steve Brooks et al. *Handbook of markov chain monte carlo*. CRC press, 2011.
- [243] Ali Javadi-Abhari et al. “Quantum computing with Qiskit”. In: *arXiv preprint arXiv:2405.08810* (2024).
- [244] Benjamin Rahn, Andrew C Doherty, and Hideo Mabuchi. “Exact performance of concatenated quantum codes”. In: *Physical Review A* 66.3 (2002), p. 032304.
- [245] Pavithran Iyer et al. “Efficient diagnostics for quantum error correction”. In: *Physical Review Research* 4.4 (2022), p. 043218.
- [246] Chris Sutherland, Todd A Brun, and Daniel A Lidar. “Non-Markovianity of the post-Markovian master equation”. In: *Physical Review A* 98.4 (2018), p. 042119.
- [247] Jalan A Ziyad, Robin Blume-Kohout, and Kenneth Rudinger. “Emergent non-Markovianity in logical qubit dynamics”. In: *arXiv preprint arXiv:2512.08893* (2025).
- [248] Antonio deMarti iOlius et al. “Decoding algorithms for surface codes”. In: *Quantum* 8 (2024), p. 1498.

- [249] J Roffe et al. “Decoding across the quantum LDPC code landscape (2020)”. In: *arXiv preprint arXiv:2005.07016* ().
- [250] Niklas Mueller et al. “Quantum computing universal thermalization dynamics in a $(2+1)$ D Lattice Gauge Theory”. In: *Nature Communications* 16.1 (2025), p. 5492.
- [251] Tzu-Chieh Wei and Paul M Goldbart. “Geometric measure of entanglement and applications to bipartite and multipartite quantum states”. In: *Physical Review A* 68.4 (2003), p. 042307.
- [252] Antonio Anna Mele. “Introduction to Haar Measure Tools in Quantum Information: A Beginner’s Tutorial”. In: *Quantum* 8 (May 2024), p. 1340. ISSN: 2521-327X. DOI: 10.22331/q-2024-05-08-1340. URL: <https://doi.org/10.22331/q-2024-05-08-1340>.
- [253] M. Ledoux. *The Concentration of Measure Phenomenon*. Mathematical surveys and monographs. American Mathematical Society, 2001. ISBN: 9780821837924. URL: <https://books.google.com/books?id=rRKChz40m6gC>.
- [254] Christian P Robert and George Casella. “Metropolis–hastings algorithms”. In: *Introducing Monte Carlo Methods with R*. Springer, 2009, pp. 167–197.
- [255] Stefanie J. Beale et al. “Quantum Error Correction Decoheres Noise”. In: *Phys. Rev. Lett.* 121 (19 Nov. 2018), p. 190501. DOI: 10.1103/PhysRevLett.121.190501. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.121.190501>.
- [256] Joseph K Iverson and John Preskill. “Coherence in logical quantum channels”. In: *New Journal of Physics* 22.7 (July 2020), p. 073066. DOI: 10.1088/1367-2630/ab8e5c. URL: <https://doi.org/10.1088/1367-2630/ab8e5c>.