Phases of matter in quantum information and error correction

by

Amirreza Negari

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Doctor of Philosophy
in
Physics

Waterloo, Ontario, Canada, 2026

© Amirreza Negari 2026

Examining Committee Membership

The following served on the Examining Committee for this thesis. The decision of the Examining Committee is by majority vote.

External Examiner:	Yi-Zhuang You Associate Professor, University of California San Diego
Supervisor(s):	Timothy Hsieh Adjunct Faculty, University of Waterloo Roger Melko Associate Professor, University of Waterloo
Internal Member:	Anton Burkov Professor, University of Waterloo Beni Yoshida Adjunct Faculty, University of Waterloo Dominic Else Adjunct Faculty, University of Waterloo
Internal-External Member:	Mohamed Hibat-Allah Assistant Professor, University of Waterloo

Author's Declaration

This thesis consists of material all of which I authored or co-authored: see Statement of Contributions included in the thesis.

This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Statement of Contributions

The chapters included in this thesis are based on collaborative work.

Chapter 2 is based on collaborative work with Subhayan Sahu and Timothy H. Hsieh [165]. Chapter 3 is based on collaborative work with Tyler D. Ellison and Timothy H. Hsieh [162]. Chapter 4 is based on collaborative work with Subhayan Sahu, Jan Behrends, Benjamin Béri, and Timothy H. Hsieh [164]. Chapter 5 is based on collaborative work with Leonardo A. Lessa and Subhayan Sahu [163]. Chapter 6 is based on collaborative work with Farzin Salek [166].

For Chapter 5, Leonardo A. Lessa and I contributed equally to the work. The project originated from numerical observations by Subhayan Sahu and me, together with my analytical explanations, which identified the central phenomenon studied in the chapter: symmetry constraints can enforce persistent high-temperature entanglement. Leonardo A. Lessa and I developed this observation into the main problem of the chapter, namely to determine when such symmetry-enforced high-temperature entanglement occurs and to characterize the structural criteria responsible for it.

My main technical contribution was the formulation and proof of the original criterion showing that symmetry constraints imply nonzero entanglement negativity at high temperature. This criterion captured the core mechanism of the chapter and is mathematically equivalent to the simplified condition used in the final presentation. Leonardo A. Lessa subsequently proved this equivalence and reformulated the criterion in a simpler and more useful form, which became the preferred formulation for later arguments. Subhayan Sahu later found a more direct derivation of this simplified condition under stronger assumptions, giving an additional conceptual and presentational simplification. These contributions together produced the equivalent formulations and technical criteria used throughout the chapter.

I also developed the first genericness argument in the discrete-symmetry setting. This argument showed that the phenomenon is not restricted to fine-tuned examples by relating the problem to local indistinguishability between canonical and Gibbs ensembles. Leonardo A. Lessa then gave a rigorous proof of the required local indistinguishability statement and developed the semiuniformity-based framework, which extended the genericness analysis to continuous symmetries such as $\mathrm{U}(1)$.

In addition, I developed the fermionic setting and showed that fermionic parity constraints realize the general mechanism studied in the chapter. Leonardo A. Lessa later generalized this perspective to arbitrary superselection rules beyond fermionic parity. He

also made substantial contributions to the proof strategy, the formalization of the arguments, and the mathematical precision of the results. All authors discussed and shaped the final presentation of the work.

Abstract

This thesis investigates phases of matter and phase transitions through the lens of quantum information, with an emphasis on phenomena not fully captured by conventional local observables or equilibrium order parameters. While the traditional framework of phase transitions relies on correlation functions and order parameters, entanglement and other information-theoretic quantities provide a broader language for characterizing both equilibrium and non-equilibrium many-body systems. A central perspective developed here is that such quantities furnish sharp diagnostics of phases and criticality, particularly in topological phases subjected to noise and measurement.

First, we study how measurements on topological quantum states reshape entanglement structure and induce phase transitions. Focusing on the toric code, we show that measuring part of the system generates distinct entanglement phases in the remaining degrees of freedom, and that tuning the measurement protocol drives transitions between them. To analyze these phenomena, we develop analytical tools that track the entanglement structure of the post-measurement state and reveal a rich phase diagram.

Next, we turn to topological codes in the presence of noise, where information-theoretic probes reveal forms of non-equilibrium criticality invisible to conventional observables. In this setting, we identify extended critical behavior in mixed states and show that conditional mutual information diagnoses transitions between distinct regimes of information retention and loss. Interpreted through quantum error correction, these transitions distinguish phases in which logical information is robustly preserved, only partially accessible, or completely lost. Building on this connection, we extend the mixed-state perspective from static codes to fault-tolerant dynamics by relating faulty syndrome-extraction circuits to the mixed-state structure of an associated higher-dimensional resource state. This leads to a decoder-independent diagnostic of fault tolerance based on the conditional mutual information of syndrome data across spacetime. The resulting spacetime Markov length diverges at the fault-tolerance threshold, providing an intrinsic information-theoretic characterization of the preservation and breakdown of logical information in noisy quantum circuits.

Finally, we develop structural results for thermal and symmetry-constrained mixed states. We show that symmetry can obstruct the sudden death of entanglement in thermal states: for canonical ensembles and for Gibbs states subject to superselection rules, entanglement persists, and in broad settings remains nonzero at arbitrarily high temperatures. In fermionic systems, this identifies parity superselection as a generic mechanism protecting mixed-state entanglement and fermionic negativity. Complementing this perspective, we

study extendibility as a tractable probe of entanglement structure in fermionic Gaussian states, showing that it admits an efficient characterization and provides practical criteria for mixed-state entanglement, including an extendibility transition in the disordered Kitaev chain.

Taken together, these results support a unified picture in which information-theoretic quantities serve as fundamental diagnostics of phase transitions and criticality in both equilibrium and non-equilibrium quantum systems.

Acknowledgements

I would like to thank my supervisor, Tim Hsieh, for his guidance, support and for always taking the time to talk.

I would also like to thank my collaborators, Subhayan Sahu, Tyler D. Ellison, Leonardo A. Lessa, Farzin Salek, Jan Behrends, Benjamin Béri, Tarun Advaith Kumar, Yijian Zou, and Roger G. Melko.

My time at PI was made much more enjoyable by the wonderful people I met here, especially Taillte. I am thankful to Leo, Jury and Jose for the many enjoyable and stimulating discussions in the sauna after squash, and to Dongjin, Subhayan, Tyler, Afshin, Amin, Barbara, Tarun, Mike, Toma, Hassan, Sercan, Yasamin, Farzin, and Zahra for many other enjoyable and stimulating discussions.

Table of Contents

Examining Committee Membership ii

Author's Declaration iii

Statement of Contributions iv

Abstract vi

Acknowledgements viii

List of Figures xv

List of Tables xxiv

1 Introduction 1

1.1 Monitored dynamics and quantum trajectories 2

1.2 Quantum error correction and syndrome monitoring 4

1.3 Error thresholds as mixed-state phase transitions 7

1.4 Mixed-state entanglement, separability, extendibility, and symmetry 9

1.5 Organization of the thesis 13

2 Measurement-induced phase transitions in the toric code **15**
2.1 Introduction . . . 15
2.2 Setup . . . 19
2.2.1 Toric code/Plaquette model . . . 19
2.2.2 Measurement setup . . . 21
2.2.3 Stabilizers and measurement . . . 22
2.3 Completely packed loop model . . . 25
2.4 Measurement-induced entanglement between two distant regions . . . 28
2.4.1 MIE between two unmeasured qubits . . . 28
2.4.2 MIE between two unmeasured boundaries . . . 29
2.5 Measurement-induced phase transition in the boundary . . . 31
2.5.1 Measured toric code as a $1+1$D hybrid circuit . . . 31
2.5.2 Bipartite entanglement in the unmeasured boundary . . . 33
2.6 Long-range order in the boundary state . . . 35
2.6.1 Spin-glass order parameter . . . 35
2.6.2 Linear order parameter from adaptive circuits . . . 36
2.7 General on-site measurements . . . 38
2.7.1 Tensor network representation of parton construction . . . 38
2.7.2 Measured toric code as a Gaussian tensor network . . . 39
2.7.3 Measured toric code as a Gaussian hybrid circuit . . . 42
2.7.4 Phase diagram of boundary MIE after general on-site measurement 44
2.8 Discussion . . . 44

3 Spacetime Markov length: a diagnostic for fault tolerance via mixed-state phases **46**
3.1 Introduction . . . 46
3.2 Mixed state phases and Markov length . . . 49
3.3 Mapping circuits to mixed states . . . 51

3.4 Mapping noisy circuits to noisy states . . . 54
3.5 Fault tolerance and mixed-state phase transition . . . 55
3.6 Classical memory in the bulk . . . 57
3.7 Spacetime Markov length . . . 61
3.8 Discussion . . . 63

4 Critical non-equilibrium phases from noisy topological memories 65
4.1 Introduction . . . 65
4.2 Setup . . . 69
4.2.1 Surface code and Majorana representation . . . 69
4.2.2 Heralded Pauli dephasing channel and the ensemble . . . 72
4.3 Phase diagram of CMI and completely packed loop model . . . 75
4.3.1 Mapping to loop model . . . 75
4.3.2 CMI computation . . . 77
4.4 Critical phase as a memory . . . 86
4.4.1 Memory and the recovery . . . 86
4.4.2 Diagnostics of global recoverability . . . 87
4.4.3 Diagnostics of local recoverability . . . 90
4.4.4 Decoding transition: global vs local . . . 96
4.5 Discussion . . . 103

5 Symmetry enforces entanglement at high temperatures 104
5.1 Introduction . . . 104
5.2 Setup and notation . . . 108
5.3 Persistence of entanglement . . . 108
5.4 Genericness and local indistinguishability . . . 109
5.5 Classification of irreps . . . 110
5.6 Persistence of negativity . . . 111
5.7 Fermions . . . 113
5.8 Discussion . . . 114

6 Extendibility of Fermionic Gaussian States **115**
6.1 Introduction . . . 115
6.2 Extendibility of fermionic states . . . 117
6.3 Extendibility of Fermionic Gaussian Channels . . . 123
6.4 Discussion . . . 124

References **126**

APPENDICES **146**

A Supplementary Material for Chapter 2 **147**
A.1 Loop patterns and quantum states in physical Hilbert space . . . 147
A.1.1 Single un-measured boundary . . . 147
A.1.2 Two un-measured boundaries . . . 149
A.2 Relation between Majorana partons and Jordan-Wigner fermions . . . 150
A.3 Jordan-Wigner transformation of the boundary after bulk measurements . 152

B Supplementary Material for Chapter 3 **155**
B.1 Feedback Channel . . . 155
B.1.1 Feedback Channel for Noiseless MBQC . . . 155
B.1.2 Feedback Channel for Noisy MBQC . . . 157
B.2 Coherent errors in the syndrome extraction circuit . . . 160
B.2.1 Coherent error and F channel . . . 160
B.2.2 Coherent error in repetition code . . . 163
B.2.3 Coherent error and forced measurements . . . 165
B.3 Coherent error in MBQC . . . 168

C Supplementary Material for Chapter 4 **171**

C.1 Punctured coherent information and local recoverability 171

C.1.1 Circuit locality and light cones 172

C.1.2 Proof of Theorem 4.1 172

C.1.3 Causal-cone insensitivity 174

C.1.4 Proof of Lemma C.1 174

C.2 Apparent critical behaviour at $q = 0$ 176

D Supplementary Material for Chapter 5 **179**

D.1 Emergence of canonical ensemble 179

D.2 Two qubit example with $\mathbb{Z}_2$ symmetry 181

D.3 Symmetric entanglement 182

D.4 Local indistinguishability of the canonical and Gibbs ensembles 183

D.5 Equivalence of SEC and EC for finite Abelian groups 189

D.6 Classification of irreps 193

D.7 Equivalence of EC and NC 197

D.8 Entanglement negativity of the thermal cluster chain state 198

D.8.1 Gibbs ensemble 199

D.8.2 Canonical ensemble 201

D.9 Persistence of fermionic negativity in the canonical ensemble 202

E Supplementary Material for Chapter 6 **203**

E.1 Preliminaries 203

E.1.1 Hilbert spaces, norms, and a convexity fact: 203

E.1.2 From qubits to fermions: modes, parity, and locality 205

E.1.3 Majorana operators and covariance matrices 206

E.1.4 Fermionic Gaussian states 207

E.1.5 Fermionic Gaussian channels 209

E.1.6 Useful lemmas: . . . 210
E.2 Extendibility of Gaussian fermions . . . 213
E.3 Proof of Theorem 6.2 . . . 215
E.4 Proof of Lemma 6.3 of main text . . . 218
E.5 Separability in Fermionic Gaussian States and Theorem 6.5 . . . 219
E.6 Proof of Theorem 6.6 . . . 222
E.7 Construction of Extendible States . . . 225
E.7.1 Beam splitter and fermionic pure–loss channel . . . 225
E.7.2 From EPR pairs to (k_1, k_2)–extendible states . . . 226
E.7.3 Bona fide condition . . . 226
E.7.4 Verification of (k_1, k_2)–extendibility . . . 227
E.8 Lower Bounds on the Distance to Fermionic Separable States . . . 228
E.8.1 Tight Lower Bound from Separable States . . . 228
E.8.2 Adapting the Bosonic Strategy . . . 230
E.9 Numerical illustration: disordered Kitaev chain . . . 231

List of Figures

2.1 (a) The bulk of a toric code ground state is measured in random Pauli bases (denoted by different colors), which induces correlations between the unmeasured 1D boundary qubits (in the dashed box). (b) Depending on the relative frequencies of different Pauli measurements, the 1D boundary can have different entanglement scaling (area law vs. logarithmic scaling) and also different orders (spin-glass vs. paramagnetic). Such phases and transitions are analyzed by mapping to a 2D classical loop model. 17

2.2 (a) Toric code stabilizers can be converted to Wen plaquette stabilizers via staggered Hadamard gates. The Wen plaquette admits a parton construction in which each qubit is split into four Majorana fermions subject to a constraint (represented by the circle in the right subfigure). (b) A ground state of the plaquette model can be constructed by projecting a free-fermion state consisting of Majorana dimers into the physical qubit Hilbert space. . 20

2.3 (a) The update process for measuring the Majorana bilinear $i\gamma_j\gamma_i$ for a free-fermion state which is a tensor product of Majorana dimers. The $+1$ refers to the measurement outcome, which sets the arrow direction in the final state. If the measurement outcome is -1 the final arrows need to be reversed. (b) Updated pairings from measuring X, Y, and Z Pauli operators on each site of the plaquette model. The updated pairings need to be tiled to form the global dimer state. We have neglected the sign tracking of the dimer pairs, which depends on the measurement outcomes. 24

2.4 The schematic phase diagram of the completely packed loop model with crossings (CPLC) [160]. The diagram distinguishes between different phases based on three order parameters, highlighted in the panel below. These order parameters correspond to loop configurations that connect points marked by red dots. The first order parameter (I) quantifies the probability of four distinct strands connecting two red points on a torus, referred to as the "watermelon correlation function." The second order parameter (II) measures the expected number of strands connecting the top and bottom boundaries on a cylinder, known as the "spanning number." The third order parameter (III) measures the expected number of strands connecting two partitions of the top boundary on a cylinder with a fixed boundary condition at the bottom. All three quantities govern measurement-induced entanglement in the toric-code state. 27

2.5 Three possible pairings of unmeasured qubits i, j after all other qubits are measured. Left: stabilizer strands prior to physical-qubit Hilbert-space projections. Right: stabilizers after projections, in canonical form. In writing the stabilizers as Pauli strings, we have assumed that the sites i, j are in the same sublattice. Configurations (a, b) are exclusively supported on i or j and do not contribute entanglement, while configuration c contributes one bit of entanglement. 29

2.6 Boundary conditions of the parton ground state of the toric code on a cylinder. The rough and smooth boundary conditions on two open edges of the cylinder are shown, along with the respective truncated stabilizers on the boundaries. The x direction is taken to have periodic boundary conditions. 30

2.7 Pauli-$X/Y/Z$ measurements on the bulk, with the boundary chain of qubits left unmeasured (left panel), correspond to a hybrid circuit with measurement and unitary gates. The measurement model is equivalent to a staggered measurement protocol on the Wen plaquette model. The plaquette model generates Majorana pairing patterns that can be interpreted as world lines of Majorana fermions undergoing a free-fermion hybrid circuit (middle panel). By Jordan-Wigner transformation, the same circuit can be identified with an Ising-symmetric measurement and unitary circuit on qubits (right panel). 32

2.8 (a) To convert any (nonlinear) spin-glass long-range order of the boundary state into (linear) ferromagnetic long-range order, a layer of on-site unitaries conditioned on bulk measurement outcomes can be applied to the boundary. (b) The classical processing for the adaptive protocol involves using the distribution of Majorana strands to construct a graph in which an edge between nodes (qubits) i, j represents the existence of a Z_iZ_j stabilizer. The edges are colored depending on whether the sign of the stabilizer is ± 1, which can be computed from the measurement outcomes along the strands. (c) The negative signs in the tree graph can be flipped by applying X_k on all nodes on one side of the negative edge (e.g. those encircled). 37

2.9 (a) The parton representation of the toric code state can be interpreted as a local tensor with one "physical" qubit leg and four "virtual" Majorana legs, along with an orientation for contraction of the virtual indices. In this case, the orientation corresponding to a toric-code ground state is specified. (b) Measuring the parton state, or equivalently, contracting the physical index with a state vector in a general direction, leads to a Gaussian state on the four "virtual" Majorana degrees of freedom. This Gaussian state can be mapped to a Gaussian operation on two Majorana fermions, and can be directly identified with a non-unitary operator. 40

2.10 Bulk measurements of general on-site operators on the toric-code ground state can be mapped to a Gaussian circuit on a chain of the virtual Majorana degrees of freedom. The qubit state supported on the top boundary on the left-hand side maps under Jordan-Wigner transformation to the Gaussian state supported on the top boundary on the right-hand side (see Appendix A.3 for derivation). 42

3.1 (a) A schematic of computing the CMI of a circuit with repeated syndrome measurements. The ensemble of measurement outcomes in spacetime is tripartitioned, and its CMI is given by $I(A:C|B) = H(AB) + H(BC) - H(B) - H(ABC)$, where $H(X)$ denotes the Shannon entropy of the bits within region X. It decays with a length scale called the spacetime Markov length ξ_p. (b) The mapping of the error threshold to a mixed-state phase transition. The circuit can be mapped to a resource state for MBQC. Measurements of the resource state decohere the state to create a mixed state. Noise in the circuit maps to noise acting on the mixed state, which can drive a mixed-state phase transition occurring at the same point as the circuit threshold. The Markov length diverges at the mixed-state phase transition, signaling a similar divergence of the spacetime Markov length at the fault-tolerance threshold. 48

3.2 Mapping the syndrome-extraction circuit of the 1D repetition code to a 2D mixed state. (a) The syndrome-extraction circuit is composed of Z_iZ_{i+1} measurements. (b) The resource state for MBQC is defined on copies of code qubits (blue) and syndrome qubits (red). The layers are labeled by integers and half-integers. There are no syndrome qubits in the $t = 0$ layer, as the input code state is assumed to be perfect. (c) The qubits are entangled into a cluster state by applying CZ gates to neighboring qubits, with stabilizers indicated. (d) Measuring the bulk in the X basis produces a mixed state coupled to the original code on the boundaries: the resulting ensemble is stabilized by local plaquette operators and composites of the logical operators of the bulk and boundary codes. 49

3.3 Error mapping from the syndrome-extraction circuit to the resource state for the 1D repetition code, as an example. Error locations are denoted by stars and are translated to the MBQC setting according to rules 1.-3. in the main text. 54

4.1 The surface code subject to heralded complete Pauli dephasing has a phase diagram described by the completely packed loop model on a square lattice, and the parameters $0 \leq p, q \leq 1$ control the relative probabilities of X, Y, and Z Paulis in the dephasing channel. There are two topologically distinct short-loop phases which exhibit finite Markov length, and an extended Goldstone phase with system-spanning loops, which exhibits diverging Markov length. We also compute the persistence of a classical logical memory using the mutual information $I(R{:}Q)$ between a classical reference bit R and the dephased code Q. In our setting, this differs from the coherent information $I_c(R\rangle Q) \coloneqq S(Q) - S(RQ)$ only by the constant $S(R) = 1$. Furthermore, we propose a punctured version of the mutual information $I(R{:}Q_{p^c})$ between the reference R and part of the code Q as a probe for local recoverability of the encoded information. Studying the interplay of phases, memory, and recoverability of encoded information in each of these phases, we find that the Goldstone phase is characterized by diverging Markov length and partial classical memory that is globally but not locally recoverable. In contrast, in the short-loop phase (shaded green), X-type logical information persists completely, and we show that it can be recovered using a quasi-local channel. 67

4.2 Left: A $d = 5$ rotated surface code with X-type (black) and Z-type (white) plaquette stabilizers. Logical operators $\bar{Z}$ (red) and $\bar{X}$ (blue) run along boundaries. Right: The Majorana–qubit tensor-network representation, with 4 virtual Majorana modes (unfilled circles) attached to each physical qubit leg (filled circles). Neighboring virtual Majoranas are projected into even-parity dimers along edges (see text), while the 4 unpaired corner modes (red circles) encode the logical qubit degrees of freedom. 70

4.3 Mapping from the dephased surface code to loop model. 74

4.4 Phase diagram for the CPLC model [160]. 77

4.5 Loop configurations and geometries for the CMI computation. 83

4.6 Numerical results for the CMI for different geometries for a fixed $q = 0.1$, with varying p and sizes of the surface code patch L, which diagnose the short loop to Goldstone phase transition in the CPLC model. 84

4.7 In the short-loop phase $p < p_c$, the decay length extracted from the global CMI (the Markov length ξ_{M}) coincides within numerical accuracy with the loop length scale ξ_{loop} extracted from the spanning number. 85

4.8 Three topologically distinct pairings of the four boundary Majorana modes induced by a trajectory, corresponding to effective logical measurement bases X, Y, and Z (left to right). Only the X-type pairing preserves the encoded classical $\bar{X}$ bit. 89

4.9 (a) Heat map of the input–output mutual information $I(R{:}Q)$ in the (p, q) plane for a surface code patch of linear size $L = 101$, computed numerically by sampling flagged measurement trajectories and identifying the induced boundary pairing sector as in Fig. 4.8. The three plateaus $I(R{:}Q) \to 1$, $1/3$, and 0 correspond to the short-loop–X, Goldstone, and short-loop–Z regimes of the loop model with crossings. Note, the grey line is a guide to the eye and not a phase boundary, rather a contour line at $I(R : Q) = 0.5$.(b) One-dimensional cut at fixed $q = 0.1$, showing $I(R{:}Q)$ versus p for several system sizes $L = 21 - 161$ with 10^4 samples, together with a finite-size scaling collapse. The crossing and collapse locate the transition at p_c and give a critical exponent $\nu \simeq 2$. 90

4.10 Punctured probes of input–output mutual information. If decoding is quasi-local, discarding a finite fraction of system size does not change the accessible information about the logical input. A persistent drop under puncturing therefore witnesses intrinsically nonlocal recovery. 93

4.11 Punctured input–output mutual information $I(R{:}Q_{p^c})$ computed using the protocol of Fig. 4.10, with a bulk square puncture of linear size $L/2$. The $I(R{:}Q) \to 1$ and $I(R{:}Q) \to 0$ short loop regimes are stable under puncturing, while the $I(R{:}Q) \to 1/3$ Goldstone regime collapse to $I(R{:}Q_{p^c}) \to 0$. The small mismatch between the apparent transition points for punctured and unpunctured curves is consistent with finite-size drift. The plot parameters are the same as Fig. 4.9. Note that the purple subregion below the transition arises from the fact that the heatmap is plotted for a fixed system size; finite size scaling reveals that the $I(R : Q_{p^c})$ goes to 1 in the $L \to \infty$ limit below the transition, as shown in (b). 94

4.12 The phase diagram from the perspective of recoverability. 97

4.13 Local loop rewirings used in the single-site update (4.27). Top row: the three possible non-canonical (Y/Z) routings of the four half-edges incident on a flagged site i (for the different choices of Pauli type (Y/Z) and sublattice), labeled (a_1)–(a_3). Bottom row: the corresponding canonical X-type routings, labeled (b_1) and (b_2). Replacing (a_i) by the appropriate (b_j) changes only the local pairing at i and therefore updates only the loop-parity constraints on the one or two strands passing through i: depending on the incident connectivity, the move either splits a strand into two or merges two strands into one, while leaving all constraints supported on loops that avoid i unchanged. 100

5.1 Schematic paths of generic many-body thermal states in relation to the set of (symmetrically) separable states SEP (SymSEP). The Gibbs ensemble $\rho_\beta \propto e^{-\beta H}$ is separable from $\beta = 0$ to β_{SDOE} due to the sudden death of entanglement, but it becomes symmetrically entangled as soon as $\beta > 0$. Similarly, the canonical ensemble $\rho_{\beta,\Lambda} \propto e^{-\beta H}\Pi_\Lambda$ is also entangled for $\beta > 0$. For a more accurate representation of SEP and SymSEP in a slice of the space of quantum states, see Fig. D.1 in Appendix D.2. 107

5.2 Entanglement negativity E_N for the Gibbs and canonical ensembles of the cluster chain Hamiltonian $H = \sum_{i=1}^{N} Z_{i-1}X_iZ_{i+1}$ under $\mathbb{Z}_2$ symmetry generated by $\prod_i X_i$ in the limit $N \to \infty$. Both curves are independent of the entangling interval size $|A| \geq 2$, and the canonical one, of the charge sector. 112

A.1 **One un-measured boundary:**(a) Stabilizer generators of the boundary state from a given pattern of loops can be obtained by multiplying a type (B) pairing (blue strand on the left) with all the interior type (A) strands (colored red in the left). **Two un-measured boundaries:** (b) Stabilizer generator obtained by multiplying two (C) type strands with all intervening (A) strands in upper and lower boundaries. (c) Two independent stabilizer generators which are products of the type (A) strands in upper and lower boundaries. 148

A.2 (a) A physical operator in terms of the parton Majoranas is represented (up to a sign) by the product of all the Majorana strands colored red in the figure. (b) The same operator can be represented as a bilinear operator of Majoranas $\tilde{\gamma}$ obtained by Jordan-Wigner transformation of the qubits. Thus, the Jordan-Wigner transformation of the qubit state is effectively the Majorana state on the $\gamma_{1,2}$ fermions. 151

A.3 The key result of this section is the cancellation between the projection tensor applied on the dimer state and the Jordan-Wigner transformation. Each large circle represents a projection tensor as defined in Chapter 2. Two of the four Majoranas per circle are paired in the depicted nearest neighbor dimer state. By contracting these tensors in this manner, we obtain a tensor that maps $2N$ Majorana fermions (bottom) to N spins (blue dots). Subsequently, the Jordan-Wigner tensor takes N spins as input and outputs $2N$ Majorana fermions. These two tensors cancel each other out when combined. 153

B.1 (a) Code qubits at times t and $t + \frac{1}{2}$ are post-selected into the $|+\rangle$ state, indicated by contraction with a triangle symbol. (b) This post-selection implements perfect teleportation, effectively enacting a time evolution from time t to $t + 1$. (c) A coherent error unitary is applied to the code qubits, represented by the blue and orange gates. The post-selection is performed afterward. The Kraus operators of both noise channels are constructed using only Pauli Z. (d) The net effect is equivalent to applying the same coherent error during teleportation. In the circuit picture, the blue channel is applied directly to the teleported qubit, while the orange channel is mapped to an equivalent channel in which each Pauli Z operator is replaced by a Pauli X. 166

C.1 Apparent critical behaviour at $q = 0.0$. Along the boundary of the CPLC phase diagram, the CMI and input–output mutual information exhibit an apparent crossing at $p^* < 1$, which we interpret as a finite-size crossover controlled by the perfect-metal fixed point at $p = 1$ rather than a true phase transition. 178

D.1 Subset of quantum states of two qubits formed by convex mixtures of the Bell states $|\Psi_-\rangle \propto |01\rangle - |10\rangle$ and $|\Phi_-\rangle = |00\rangle - |11\rangle$, and the maximally mixed $X_1X_2 = +1$ state $\frac{1}{2}\Pi_+ = \frac{1}{2}|00\rangle\langle 00| + \frac{1}{2}|11\rangle\langle 11|$. The dashed line traces the path of the Heisenberg antiferromagnetic thermal states $e^{-\beta \mathbf{S}_1 \cdot \mathbf{S}_2}$, or, equivalently, of the Werner states $\lambda |\Psi_-\rangle\langle\Psi_-| + (1 - \lambda)\frac{1}{4}\mathbb{1}$. Meanwhile, the thick black line going from $\frac{1}{2}\Pi_-$ to $|\Psi_-\rangle$ traces the path of the canonical ensemble $e^{-\beta \mathbf{S}_1 \cdot \mathbf{S}_2}\Pi_{-1}$ in the $X_1X_2 = -1$ sector. The SEP region in cyan consists of separable states, with the blue SymSEP line inside it indicating the symmetrically separable ones. Their loci were found via the Peres-Horodecki criterion, since it is necessary and sufficient for two-qubit systems [171, 94]. This illustration was inspired by Fig. 16.8(b) of [19]. 182

E.1 **Disordered Kitaev chain: probability of two-site (k_1, k_2)-extendibility versus disorder.** For the disordered Kitaev chain (E.97) with periodic boundary conditions and i.i.d. chemical-potential disorder $\mu_j \sim \mathrm{Unif}[-W, W]$, we compute the Gaussian ground-state covariance matrix, reduce it to two adjacent sites, and test two-sided (k_1, k_2)-extendibility of the resulting two-site Gaussian state using the covariance-matrix SDP characterization from the main text. The plotted quantity is the empirical probability $p(W)$ over disorder realizations. The rise of $p(W)$ with W indicates that stronger disorder makes the local reduced state more likely to satisfy the extendibility constraints. Parameters: $N = 10$, $(k_1, k_2) = (1, 2)$, and $N_{\mathrm{samp}} = 1000$. . . . 233

List of Tables

2.1 Circuit mapping of the (p, q) measurement model on the toric code (TC). For completeness, the corresponding staggered measurement protocol for the Wen plaquette (WP) model is mentioned in parentheses. 34

5.1 Typical behavior of entanglement measures (E, E_S) as functions of the inverse temperature $\beta = T^{-1}$ for the Gibbs and canonical ensembles in the cases with and without superselection. E (E_S) is a measure of (symmetric) entanglement, such as (fermionic) negativity. There is sudden death of entanglement for the Gibbs ensemble with no superselection rule if H is local [12], while there is not for the other cases if their respective entangling conditions (EC and SEC) are satisfied. For an example of the above, see Fig. 5.2. 105

Chapter 1

Introduction

The theory of phases of matter is one of the central organizing principles of many-body physics. Traditionally, phases and phase transitions have been characterized in terms of local observables, correlation functions, and symmetry breaking. This framework successfully describes a wide range of equilibrium phenomena, but it is incomplete for many quantum many-body systems of current interest in condensed matter physics and quantum information. In particular, topological phases, monitored quantum systems, noisy quantum memories, and open quantum systems exhibit structures that are not fully captured by conventional order parameters [232, 112, 140, 211, 14, 53, 130].

These systems share a common feature: their physically relevant properties are often encoded not only in local expectation values, but also in the structure of quantum information. Quantities such as entanglement, conditional mutual information, and coherent information therefore provide a broader language for characterizing phases and phase transitions. In monitored systems, measurements generate ensembles of quantum trajectories whose entanglement structure can itself exhibit distinct phases. In noisy topological codes, information-theoretic quantities diagnose transitions associated with the preservation or loss of logical information. Even in equilibrium mixed states, there are transitions such as the sudden death of entanglement, which can be understood as information-theoretic transitions and are invisible to conventional linear order parameters [140, 211, 68, 193, 12].

This thesis studies these questions within a unified information-theoretic perspective on quantum many-body physics. In Chapters 2–4, I investigate topological and error-correcting systems subjected to measurements and noise, focusing on measurement-induced phases, mixed-state criticality, and fault tolerance. In Chapters 5 and 6, I turn to structural questions about mixed-state entanglement, studying the role of symmetry in thermal states

and the use of extendibility in fermionic Gaussian systems. I first review the necessary background and conceptual framework for these topics in order to set the stage for the later chapters.

1.1 Monitored dynamics and quantum trajectories

A monitored quantum system is naturally described at the level of individual measurement records. Starting from an initial pure state $|G\rangle$, a sequence of measurements produces an ensemble of conditional states, or *quantum trajectories*. For a realization r of measurement choices and outcomes, the post-measurement state is

$$|\psi_r\rangle = \frac{K_r|G\rangle}{\sqrt{p_r}}, \qquad p_r = \langle G|K_r^\dagger K_r|G\rangle, \tag{1.1}$$

where K_r is the Kraus operator associated with the record r. The Kraus operators satisfy the completeness relation

$$\sum_r K_r^\dagger K_r = \mathbb{I}. \tag{1.2}$$

If the measurement outcomes are discarded, one instead obtains the averaged state

$$\bar{\rho} = \sum_r p_r |\psi_r\rangle\langle\psi_r|. \tag{1.3}$$

Thus, monitored evolution gives rise to two distinct objects: an ensemble of pure conditional states and an unconditional mixed state obtained by averaging over the measurement record.

For the purposes of this thesis, the essential distinction is between quantities that are linear and nonlinear in this ensemble. If $\mathcal{O}$ is linear in the density matrix, then

$$\sum_r p_r \operatorname{Tr}(\mathcal{O}\rho_r) = \operatorname{Tr}(\mathcal{O}\bar{\rho}), \tag{1.4}$$

so averaging over trajectories is equivalent to evaluating the observable in the mixed state $\bar{\rho}$. By contrast, entanglement is nonlinear. For a bipartition $A \cup \bar{A}$, the entanglement entropy of trajectory r is

$$S_A^{(r)} = -\operatorname{Tr}\big(\rho_A^{(r)} \log \rho_A^{(r)}\big), \qquad \rho_A^{(r)} = \operatorname{Tr}_{\bar{A}} |\psi_r\rangle\langle\psi_r|, \tag{1.5}$$

and in general

$$\sum_r p_r S_A^{(r)} \neq S\left(\sum_r p_r \rho_A^{(r)}\right). \tag{1.6}$$

This is the basic origin of measurement-induced entanglement: a monitored system can undergo sharp changes in trajectory-averaged entanglement even when no singular behavior is visible in linear observables of the averaged state [140, 211, 141, 79, 239, 142].

A particularly useful setting is one in which only a subset M of the system is measured, while the complement M^c remains unmeasured. For each realization r, the measurements on M prepare a conditional state on M^c,

$$\rho_{M^c}^{(r)} = \mathrm{Tr}_M \, |\psi_r\rangle\langle\psi_r|. \tag{1.7}$$

For the projective measurement protocols considered in this thesis, $\rho_{M^c}^{(r)}$ is itself pure for each realization. The physically relevant object is therefore not a single state on M^c, but an ensemble $\{p_r, \rho_{M^c}^{(r)}\}$. One may then ask how the entanglement of the conditional states scales with subsystem size and geometry, and whether tuning the measurement protocol can drive transitions between distinct entanglement patterns.

This viewpoint is particularly natural for entangled resource states. Measurements on part of the system do not merely remove degrees of freedom; rather, they induce a nontrivial ensemble of states on the unmeasured qubits. In this sense, monitored dynamics acts as a state-preparation protocol on M^c, with the measurement record selecting a particular trajectory within the ensemble. This is closely related to the logic of measurement-based quantum computation, where measurements on a higher-dimensional entangled state effectively generate dynamics or state preparation on a lower-dimensional boundary [180].

Chapter 2 develops this picture in a concrete topological setting. There, random single-qubit measurements in the bulk of the toric code generate distinct phases on an unmeasured one-dimensional boundary. Varying the probabilities of different Pauli measurements drives transitions between phases with different entanglement scaling and different order parameters. A key advantage of this setup is that the induced boundary states can be analyzed analytically by mapping the post-measurement ensemble to Gaussian tensor networks and, in the Pauli case, to a classical loop model [165]. In this way, the chapter shows how measurements on a topological resource state can be used to generate and diagnose distinct phases of matter.

1.2 Quantum error correction and syndrome monitoring

A second setting in which measurement records play a central role is quantum error correction. The basic idea is to store logical information nonlocally, so that local noise does not immediately destroy it [209]. In a stabilizer code, the logical Hilbert space is embedded into a larger physical Hilbert space and defined as the simultaneous $+1$ eigenspace of a commuting family of Pauli operators $\{S_a\}$:

$$\mathcal{H}_{\text{code}} = \{|\psi\rangle : S_a|\psi\rangle = |\psi\rangle,\ \forall a\}\,. \tag{1.8}$$

Logical operators commute with all stabilizers but act nontrivially within $\mathcal{H}_{\text{code}}$. The essential point is that one does not measure the logical state directly. Instead, one measures the stabilizers, whose outcomes reveal the presence of errors while preserving the encoded logical degrees of freedom.

A canonical example is the toric code, in which qubits live on the edges of a square lattice and the stabilizers are

$$A_v = \prod_{e \ni v} X_e, \qquad B_p = \prod_{e \in \partial p} Z_e. \tag{1.9}$$

The code space is the simultaneous $+1$ eigenspace of all A_v and B_p [112]. A local Pauli error creates a characteristic pattern of violated stabilizers: a single Z error flips the eigenvalues of adjacent star operators, while a single X error flips adjacent plaquette operators. More generally, if an error operator E anticommutes with a stabilizer S_a, then the outcome of measuring S_a changes sign. The collection of all such outcomes is the *syndrome*. Since many microscopic errors differ only by stabilizers and therefore act identically on the logical information, the goal of decoding is not to identify the exact physical error, but rather to infer an equivalent error class and apply a correction with the same syndrome and opposite logical action.

If $\mathcal{N}$ denotes a noise channel and $\mathcal{M}_s$ denotes projection onto syndrome sector s, then a single round of syndrome monitoring takes the schematic form

$$\rho \mapsto \rho_s = \frac{\mathcal{M}_s \circ \mathcal{N}(\rho)}{\mathrm{Tr}[\mathcal{M}_s \circ \mathcal{N}(\rho)]}, \qquad p_s = \mathrm{Tr}[\mathcal{M}_s \circ \mathcal{N}(\rho)]. \tag{1.10}$$

Thus, syndrome extraction is itself a monitored process: the syndrome is a classical measurement record, and each outcome prepares a conditional post-measurement state. In this sense, quantum error correction naturally fits into the broader language of monitored

dynamics introduced above. For the present discussion, it is enough to focus on the simpler setting in which recovery is based on the syndrome information produced by noise acting on the code.

From the perspective of quantum information, the central question is whether the effect of the noise can be *undone*. Suppose a logical state ρ_L is encoded into a physical code state ρ_{code}, and is then subjected to a noise channel $\mathcal{N}$. A recovery map $\mathcal{R}$ corrects the noise if the composition $\mathcal{R} \circ \mathcal{N}$ acts approximately as the identity on the code subspace:

$$(\mathcal{R} \circ \mathcal{N})(\rho_{\text{code}}) \approx \rho_{\text{code}} \qquad \text{for all encoded states } \rho_{\text{code}}. \tag{1.11}$$

In this sense, quantum error correction is the problem of approximately inverting a noisy channel on the encoded degrees of freedom. The syndrome record provides partial information about the error, and the decoder uses that information to construct a recovery map. When such an inversion is possible with arbitrarily high accuracy as the code size increases, the encoded information is robust; when it is not, the logical information is eventually lost.

A natural way to formalize this notion of recoverability is to entangle the logical input with a reference system. Let L denote the logical subsystem and R an isomorphic reference, initially prepared in a maximally entangled state

$$|\Phi\rangle_{RL} = \frac{1}{\sqrt{d_L}} \sum_{\alpha=1}^{d_L} |\alpha\rangle_R \otimes |\alpha\rangle_L, \tag{1.12}$$

where d_L is the logical Hilbert-space dimension. After encoding and applying noise, but before any recovery is performed, the joint state of the reference and output is

$$\rho_{RQ}^{\text{noisy}} = (\text{id}_R \otimes \mathcal{N}_{d,p})\big(|\Phi\rangle\langle\Phi|_{RL}\big), \tag{1.13}$$

where $\mathcal{N}_{d,p}$ denotes the effective noisy channel for a code of distance d at physical noise strength p. If the encoded information survives the noise, then the output remains almost maximally entangled with the reference; if the logical information is lost, this entanglement is degraded. The advantage of this formulation is that it allows one to ask about recoverability directly at the level of the noisy state, before specifying any particular decoder.

This motivates the use of the *coherent information*,

$$I_c(R\rangle Q)_\rho = -S(R|Q)_\rho = S(Q)_\rho - S(RQ)_\rho. \tag{1.14}$$

For a maximally entangled state of Schmidt rank d_L, the coherent information is maximal:

$$I_c(R\rangle Q) = \log d_L. \tag{1.15}$$

Thus, from the error-correction perspective, the coherent information of the noisy state measures how much of the logical quantum information remains correlated with the reference and is therefore still available, in principle, to be recovered.

This language is useful because coherent information cannot increase under further processing of the output. In particular, if one now applies a recovery channel $\mathcal{R}$ to the noisy state, then

$$I_c(R\rangle Q)_{(\mathrm{id}_R \otimes (\mathcal{R} \circ \mathcal{N}_{d,p}))(\Phi_{RL})} \leq I_c(R\rangle Q)_{\rho_{RQ}^{\mathrm{noisy}}}. \tag{1.16}$$

Therefore, successful recovery requires that the noisy state itself already retain essentially all of its coherent information. If there exists a sequence of recovery maps $\mathcal{R}_d$ such that $(\mathcal{R}_d \circ \mathcal{N}_{d,p})$ approaches the identity channel on the encoded subspace as $d \to \infty$, then necessarily

$$I_c(R\rangle Q)_{\rho_{RQ}^{\mathrm{noisy}}} \to \log d_L. \tag{1.17}$$

Conversely, if the coherent information of the noisy state is asymptotically close to $\log d_L$, then the noisy channel preserves essentially all entanglement with the reference, which means that the action of the noise on the encoded subspace is asymptotically reversible. In this way, near-maximal coherent information of the noisy state provides a natural information-theoretic criterion for recoverability.

This point of view gives a natural formulation of the error-correction threshold. Physically, the threshold is the largest noise strength below which increasing the redundancy of the code continues to improve its ability to reverse the effect of local noise on the encoded state. Below threshold, typical errors remain sparse enough that the syndrome contains sufficient information to identify an appropriate correction, and increasing the code distance suppresses logical failure. Above threshold, errors proliferate to the point that distinct logical sectors can no longer be reliably distinguished, and the encoded information is no longer fully recoverable in the thermodynamic limit.

The *fundamental* threshold of the code is obtained only after optimizing over all admissible recovery strategies. Equivalently, it is the largest p for which there exists some sequence of recovery maps achieving asymptotically perfect correction. Because the recovered coherent information can never exceed the coherent information already present in the noisy state, this optimal threshold can be characterized directly from the noisy channel:

$$p_{\mathrm{th}} = \sup \left\{ p : \; I_c(R\rangle Q)_{\rho_{RQ}^{\mathrm{noisy}}} \to \log d_L \text{ as } d \to \infty \right\}. \tag{1.18}$$

Below this optimal threshold, the noisy state still retains essentially all of its entanglement with the reference, so some sufficiently good decoder can asymptotically recover the full logical state. Above threshold, no decoder can do so. Importantly, this does *not* mean that

the coherent information must vanish above threshold. Rather, it means that the coherent information is asymptotically strictly smaller than its maximal value, so full recovery of the encoded state is no longer possible, even though some logical information may still survive.

In this way, the threshold may be understood as the point where the noisy channel ceases to be asymptotically invertible on the encoded degrees of freedom. This recoverability-based viewpoint will be important in the following chapters, where the survival, loss, and locality properties of encoded information are related to mixed-state phases and information-theoretic diagnostics.

1.3 Error thresholds as mixed-state phase transitions

The discussion above suggests that the threshold of a quantum memory should not be viewed only as a property of a decoder. Rather, it is natural to ask whether the threshold reflects a more intrinsic change in the structure of the noisy state itself. For topological codes, this question leads directly to the modern notion of a *mixed-state phase*.

For pure ground states, two states are said to be in the same phase if they can be connected by finite-time evolution generated by a local Hamiltonian, or equivalently by a finite-depth local unitary circuit. For mixed states, the natural analogue is local dissipative evolution. A standard definition therefore says that two mixed states ρ_1 and ρ_2 are in the same phase if there exist local Lindbladian evolutions that connect them in both directions [53]:

$$\mathcal{T}e^{\int_0^1 dt\,\mathcal{L}_1(t)}[\rho_1] = \rho_2, \qquad \mathcal{T}e^{\int_0^1 dt\,\mathcal{L}_2(t)}[\rho_2] = \rho_1. \tag{1.19}$$

This two-way condition is the mixed-state analogue of finite-depth local-unitary equivalence. The physical idea is simple: a local channel can add noise, but if the final state can still be returned to the initial one by another local channel, then no genuinely new long-range structure has been created or destroyed.

While this two-way notion is natural, it is not by itself sufficiently sensitive for the purposes of quantum memories. In the context of error correction, what matters is not only that a noisy state can in principle be connected back to the original state, but that the noise can be *locally inverted*. This leads to the sharper notion of *local reversibility*. The key quantity controlling local reversibility is the conditional mutual information (CMI) [193]. For a tripartite state ρ_{ABC}, the CMI is defined by

$$I(A:C|B)_\rho = S(AB) + S(BC) - S(B) - S(ABC), \tag{1.20}$$

where $S(X) = -\operatorname{Tr}(\rho_X \log \rho_X)$ is the von Neumann entropy. If A and C are separated by a buffer region B of linear size r, then the decay of the CMI with r defines a characteristic length scale, the *Markov length* ξ_{M}, through

$$I(A:C|B)_\rho \sim e^{-r/\xi_{\mathrm{M}}}. \tag{1.21}$$

Physically, finite Markov length means that the correlations between A and C are screened by a finite buffer and the effect of local noise can be reversed quasi-locally, whereas a diverging Markov length signals a breakdown of such local recoverability.

A particularly instructive example is the toric code subject to local bit-flip noise. Let $|\mathrm{t.c.}\rangle$ denote a ground state of the toric code and consider the noisy state

$$\rho_p = \mathcal{E}_p^x\left[|\mathrm{t.c.}\rangle\langle\mathrm{t.c.}|\right], \tag{1.22}$$

where

$$\mathcal{E}_p^x[\cdot] = \bigotimes_i \left((1-p)\ \cdot + p\, X_i(\cdot)X_i\right). \tag{1.23}$$

As the bit-flip rate p is increased, the logical $\tilde{Z}$ information eventually becomes undecodable, while the logical $\tilde{X}$ information survives because it commutes with the noise. Operationally, this is the usual decodability threshold of the toric code. From the mixed-state perspective, however, one asks a different question: can the noisy state ρ_p still be locally reversed to the original topological memory?

The answer is encoded in the Markov length. For the noisy toric-code state, the CMI can be related to the free-energy cost of point defects in the corresponding random-bond Ising model, and this makes it possible to compute the Markov length analytically [193]. The result is that the Markov length remains finite away from the threshold and diverges precisely at the decodability transition. Equivalently, the threshold is the point at which the noisy state ceases to be quasi-locally reversible to the original code state. Below threshold, the effect of the bit-flip noise can still be undone by a quasi-local recovery channel; above threshold, no such quasi-local inversion of the noise is possible. In this way, the ordinary error-correction threshold is reinterpreted as a mixed-state phase transition diagnosed by CMI [68, 193, 14].

This identification is conceptually significant. In equilibrium critical phenomena, one usually associates a phase transition with a diverging correlation length of local observables. Here, by contrast, the singular quantity is not an ordinary correlation length but the Markov length extracted from conditional mutual information. The transition is therefore

invisible to conventional local order parameters, yet it is sharply visible in the information-theoretic structure of the state. What becomes critical is not a local observable, but the ability to locally reconstruct the encoded quantum information.

This perspective will be important in the following chapters. Chapter 3 extends the same logic to repeated syndrome extraction and fault-tolerant dynamics, where one defines an analogous Markov length directly in the spacetime distribution of syndrome outcomes. Chapter 4 goes further and studies a setting in which the Markov length does not merely diverge at an isolated threshold, but exhibits divergence throughout an extended non-equilibrium critical phase.

1.4 Mixed-state entanglement, separability, extendibility, and symmetry

The latter part of this thesis turns from noisy topological systems to structural questions about mixed-state entanglement. For pure states, entanglement is completely characterized by the impossibility of writing the state as a tensor product across a bipartition. For mixed states, the situation is subtler. A mixed state may contain both classical and quantum correlations, and one must therefore distinguish genuine entanglement from statistical mixing.

For a bipartite mixed state ρ_{AB}, the basic notion is *separability*. One says that ρ_{AB} is separable if it can be written as

$$\rho_{AB} = \sum_{\alpha} p_\alpha \, \rho_A^{(\alpha)} \otimes \rho_B^{(\alpha)}, \qquad p_\alpha \geq 0, \qquad \sum_{\alpha} p_\alpha = 1. \tag{1.24}$$

If no such decomposition exists, then ρ_{AB} is entangled [234]. This definition already illustrates why mixed-state entanglement is more difficult than its pure-state counterpart: one must optimize over all possible decompositions of the density matrix, and in general this is a hard problem. Moreover, unlike for pure states, the entropy of a subsystem is no longer by itself an entanglement measure, since it also contains contributions from classical mixing.

For this reason, one often works with entanglement measures and diagnostics tailored to mixed states. One important example is the *entanglement of formation*, defined as

$$E_F(\rho_{AB}) = \inf_{\{p_\alpha, |\psi_\alpha\rangle\}} \sum_{\alpha} p_\alpha \, S(\mathrm{Tr}_B \, |\psi_\alpha\rangle\langle\psi_\alpha|) \,, \tag{1.25}$$

where the infimum is taken over all pure-state decompositions

$$\rho_{AB} = \sum_{\alpha} p_\alpha |\psi_\alpha\rangle\langle\psi_\alpha|. \tag{1.26}$$

This quantity has a clear operational meaning, but it is generally difficult to compute because of the optimization over all decompositions. Other widely used measures include the relative entropy of entanglement and squashed entanglement, which are also conceptually powerful but often hard to evaluate in practice.

A more tractable example is the entanglement negativity. Given a bipartite state ρ_{AB}, one defines its partial transpose with respect to subsystem B, denoted $\rho_{AB}^{T_B}$, and from it the negativity

$$\mathcal{N}(\rho_{AB}) = \frac{\|\rho_{AB}^{T_B}\|_1 - 1}{2}, \tag{1.27}$$

or equivalently the logarithmic negativity

$$E_{\mathcal{N}}(\rho_{AB}) = \log \|\rho_{AB}^{T_B}\|_1. \tag{1.28}$$

These quantities vanish for separable states and are often computationally accessible, which makes them useful probes of whether entanglement survives in thermal or noisy mixed states [171, 94, 223]. At the same time, negativity is only a sufficient diagnostic: there exist entangled mixed states with vanishing negativity, so no single computable witness completely characterizes mixed-state entanglement.

A complementary approach is provided by *extendibility*, which gives a hierarchy of increasingly stringent approximations to separability. A bipartite state ρ^{AB} is said to be (k_1, k_2)-extendible if there exists a larger state

$$\widetilde{\rho}^{A_1\cdots A_{k_1} B_1 \cdots B_{k_2}} \tag{1.29}$$

that is invariant under permutations of the A systems and, separately, of the B systems, and whose two-body marginals satisfy

$$\widetilde{\rho}^{A_i B_j} = \rho^{AB} \qquad \forall\, i, j. \tag{1.30}$$

Intuitively, extendibility quantifies how freely the correlations in ρ^{AB} can be shared with additional copies of either subsystem. Since entanglement is monogamous, strongly entangled states cannot be extended to many copies in this way. In the limit of infinite extendibility, one recovers separability: a state is separable if and only if it admits symmetric extensions

to arbitrarily many copies [59, 60, 138, 69, 61, 219, 48]. Finite extendibility therefore provides a controlled approximation to the separable set and plays the role of a witness to entanglement.

For mixed states, symmetry introduces an additional layer of structure. Unlike for pure states, symmetry can constrain not only the density matrix itself but also the admissible decompositions that enter the definition of separability. Let G be a symmetry group represented on the Hilbert space by unitaries $U(g)$. A mixed state ρ is said to be *weakly symmetric* if

$$U(g)\rho U(g)^\dagger = \rho \qquad \forall g \in G. \tag{1.31}$$

For a continuous $U(1)$ symmetry generated by a conserved charge Q, this is equivalently

$$[\rho, Q] = 0. \tag{1.32}$$

Weak symmetry means that the state is invariant under the symmetry after averaging, but it does not constrain the individual pure or mixed states that may appear in a decomposition of ρ [38, 7, 6].

A stronger notion is obtained by imposing symmetry directly at the level of the ensemble. Given a one-dimensional irrep or symmetry charge $\Lambda : G \to U(1)$, one may restrict the ensemble to states $|\psi\rangle$ satisfying

$$U(g)|\psi\rangle = \Lambda(g)|\psi\rangle \qquad \forall g \in G. \tag{1.33}$$

A mixed state built entirely from such states is said to be *strongly symmetric*. Equivalently, if Π_Λ denotes the projector onto the corresponding charge sector, then a strongly symmetric state satisfies

$$U(g)\rho_\Lambda = \Lambda(g)\rho_\Lambda, \qquad \text{or equivalently} \qquad \rho_\Lambda = \Pi_\Lambda \rho_\Lambda \Pi_\Lambda. \tag{1.34}$$

Thus weak symmetry means invariance under the symmetry, whereas strong symmetry means that the state is supported entirely within a definite symmetry sector.

The relation between weak and strong symmetry can be understood through the decomposition of the Hilbert space into symmetry sectors. For Abelian G, the Hilbert space decomposes into isotypic components labeled by irreducible representations λ, and any weakly symmetric state can be written as a convex combination of strongly symmetric blocks,

$$\rho = \sum_\lambda q_\lambda \rho_\lambda, \qquad \rho_\lambda = \frac{\rho \Pi_\lambda}{q_\lambda}, \qquad q_\lambda = \mathrm{Tr}(\rho \Pi_\lambda). \tag{1.35}$$

In this sense, weak symmetry may be viewed as a mixture over strongly symmetric sectors, whereas strong symmetry corresponds to selecting a single such sector. For mixed states, this distinction is important because it changes which decompositions are available in defining separability.

A related but conceptually distinct idea is that of a *superselection rule*. Here one does not merely impose symmetry on a particular state; rather, one restricts the quantum theory itself to states and operations compatible with the symmetry. In a theory with superselection, only weakly symmetric states are regarded as physical. This changes not only the dynamics but also the notion of entanglement itself.

More precisely, a bipartite state ρ_{AB} is called *symmetrically separable* if it admits a decomposition

$$\rho_{AB} = \sum_i p_i \, \rho_{A,i} \otimes \rho_{B,i}, \qquad p_i \geq 0, \qquad \sum_i p_i = 1, \tag{1.36}$$

such that each local factor $\rho_{A,i}$ and $\rho_{B,i}$ is itself a valid state within the superselected theory, and therefore weakly symmetric on its own. This requirement is generally stronger than ordinary separability: there can exist states that are separable in the usual sense, but not symmetrically separable once the symmetry constraint is imposed on the local components [15, 74, 40, 221, 199, 201, 16, 166, 234].

The example most relevant for this thesis is fermion parity superselection. If $P = (-1)^{\hat{N}}$ denotes fermion parity, then physical states must commute with P,

$$[\rho, P] = 0. \tag{1.37}$$

This forbids coherent superpositions of states with even and odd fermion number. For instance, for a single fermionic mode, the pure state

$$|\psi\rangle = \frac{1}{\sqrt{2}}\big(|0\rangle + |1\rangle\big) \tag{1.38}$$

is not a physical state in the presence of parity superselection, because it does not have definite parity and the corresponding density matrix does not commute with P. By contrast, the incoherent mixture

$$\rho = \frac{1}{2}|0\rangle\langle 0| + \frac{1}{2}|1\rangle\langle 1| \tag{1.39}$$

is physical, since it is parity-even as a density matrix. Similarly, for a bipartite fermionic system, a decomposition into local product states is only allowed if each local factor is itself parity symmetric. This is why superselection can turn a state that is separable in

the ordinary sense into one that is entangled once separability is evaluated within the symmetry-restricted theory [91, 16, 166].

In this way, superselection provides a second route by which symmetry reshapes mixed-state entanglement. Whereas strong symmetry restricts the ensemble to a fixed charge sector, a superselection rule restricts the set of allowed local states and thereby changes the geometry of the separable set itself. As a result, a state that appears unentangled in the usual quantum theory may become entangled once separability is evaluated within the symmetry-restricted theory.

These observations are important for understanding mixed-state phases and transitions. Even in equilibrium, mixed states can undergo information-theoretic transitions that are invisible to conventional linear order parameters. A notable example is the sudden death of entanglement in thermal states, where increasing temperature drives a transition from an entangled mixed state to one that is exactly separable [12]. Chapter 5 studies how symmetry constrains mixed-state entanglement in thermal ensembles, while Chapter 6 develops extendibility as a quantitative probe of entanglement structure. The broader message is that, for mixed states, entanglement is not a single binary property but part of a richer geometric structure involving separability, hierarchy approximations such as extendibility, and symmetry constraints on the space of allowed states.

1.5 Organization of the thesis

The remainder of this thesis is organized as follows.

Chapter 2 studies measurement-induced phases generated by random local measurements on the toric code. It shows how measurements on the bulk induce distinct phases on an unmeasured boundary, characterized by different entanglement scaling and order.

Chapter 3 develops a mixed-state interpretation of fault tolerance. It introduces the spacetime Markov length as a decoder-independent diagnostic of the threshold by relating noisy syndrome-extraction circuits to decohered resource states in one higher dimension.

Chapter 4 investigates noisy topological memories and identifies an extended non-equilibrium critical phase with divergent Markov length. It also distinguishes between globally recoverable logical information and information accessible to quasi-local decoders.

Chapter 5 studies thermal mixed states constrained by symmetry. It shows that symmetry and superselection can prevent the sudden death of entanglement and allow entanglement to persist at arbitrarily high temperatures in broad settings.

Chapter 6 studies extendibility as a probe of mixed-state entanglement in fermionic Gaussian systems. It develops an efficient characterization of extendibility and explores its consequences for separability and channel structure.

Taken together, these chapters support a unified perspective in which information-theoretic quantities serve as fundamental probes of phases, transitions in both equilibrium and non-equilibrium quantum systems.

Chapter 2

Measurement-induced phase transitions in the toric code

This chapter shows how distinct phases of matter can be generated by performing random single-qubit measurements on a subsystem of the toric code. Using a parton construction, such measurements map to random Gaussian tensor networks, and in particular, random Pauli measurements map to a classical loop model in which watermelon correlators precisely determine measurement-induced entanglement. Measuring all but a one-dimensional boundary of qubits realizes hybrid circuits involving unitary gates and projective measurements in $1+1$ dimensions. We find that varying the probabilities of different Pauli measurements can drive transitions in the unmeasured boundary between phases with different orders and entanglement scaling, corresponding to short- and long-loop phases in the classical model. Furthermore, by utilizing single-site boundary unitaries conditioned on the bulk measurement outcomes, we generate mixed-state ordered phases and transitions that can be experimentally diagnosed via linear observables. This demonstrates how parton constructions provide a natural framework for measurement-based quantum computing setups to produce and manipulate phases of matter.

2.1 Introduction

Investigating the quantum phases of matter that can be dynamically generated in a quantum processor using measurement, classical feedback, and local unitaries has been a fruitful area of research. There has been significant interest in using such hybrid circuits to manipulate entanglement patterns, starting with the observation of a measurement-induced

entanglement transition from volume-law to area-law scaling as the frequency of measurements is varied [2, 140, 211, 141, 79, 239, 142, 103, 13] (for a review, see [71, 175]). Even without any unitary gates, random measurements of multi-site operators can lead not only to various entanglement patterns but also to distinct long-range orders, such as symmetry-protected topological order, spin-glass order, and intrinsic topological order [192, 125, 126, 214]. These orders can undergo phase transitions by adjusting the probabilities of competing measurements.

A different context in which measurements take center stage is measurement-based quantum computation (MBQC) [180]. The MBQC approach involves starting with an entangled "resource state," such as the 2D cluster state, and sequentially performing single-site measurements on the majority of the qubits, where the measurement basis can depend on the outcomes of previous measurements. This results in the remaining unmeasured qubits being directed towards a specific entangled state that encodes the outcome of a deterministic quantum computation. For example, measurements on a 2D resource state effectively realize a computation on the 1D boundary of the system, and the other dimension corresponds to the "time" direction of the computation.

In this chapter, we ask the following question: starting with an entangled resource state, can MBQC-type protocols lead to robust quantum phases of matter and transitions between them? We explore this question for the toric code ground state, which is an exactly solvable model of $\mathbb{Z}_2$ topological order in two dimensions [110]. We find that by tuning the probabilities of measuring single-site Pauli X, Y, or Z in the toric code bulk, we can realize distinct phases in an unmeasured boundary. These measurement-induced phases are characterized by the presence or absence of spin-glass order parameters and by their entanglement scaling (area law vs. logarithmic scaling). As is the case for MBQC, the bulk measurements in the toric code effectively realize dynamics for a one-dimensional boundary, and in this case the effective dynamics are those of a hybrid circuit involving both unitaries and measurements. Thanks to the underlying entanglement of the toric code state, only single-qubit measurements are required to effectively realize nontrivial hybrid circuits involving two-qubit operations.

Liu et al. [147] performed a related study on the 2D cluster state, an MBQC resource state which enables universal quantum computation, and discovered an entanglement transition from area to volume law in boundary qubits induced by measuring the bulk qubits. In contrast, in our setup with toric code, we find transitions in both entanglement (albeit without volume law) and other order parameters. A recent study [80] also considered an MBQC setup on the 2D cluster state and found evidence of distinct area-law entanglement phases on the 1D boundary. One advantage of our setup is that we can understand such transitions analytically by relating the entanglement properties of the qubits to certain

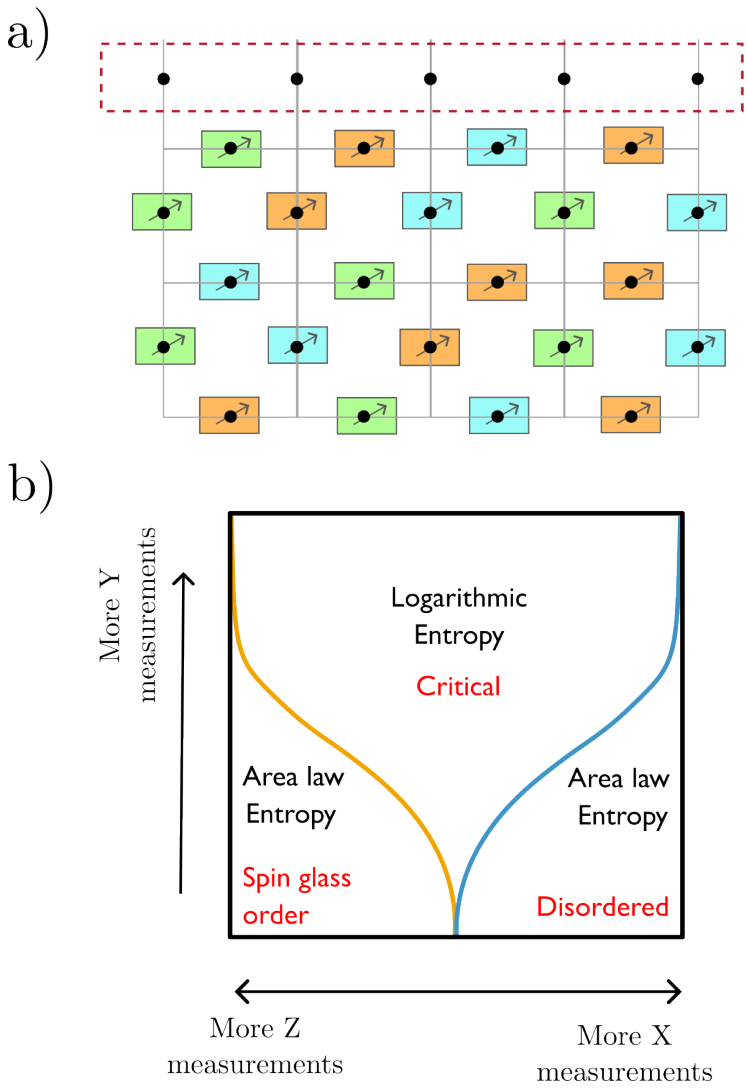


Figure 2.1: (a) The bulk of a toric code ground state is measured in random Pauli bases (denoted by different colors), which induces correlations between the unmeasured 1D boundary qubits (in the dashed box). (b) Depending on the relative frequencies of different Pauli measurements, the 1D boundary can have different entanglement scaling (area law vs. logarithmic scaling) and also different orders (spin-glass vs. paramagnetic). Such phases and transitions are analyzed by mapping to a 2D classical loop model.

correlation functions of a corresponding 2D classical loop model with crossings. Such a model has short- and long-loop phases, which exactly correspond to the area law and the logarithmic scaling of entanglement in the 1D boundary. A summary of these results is presented in Fig. 2.1. We also note that [17] established a very different mapping between toric code error correction in the presence of both incoherent and coherent errors and $(1+1)$D free-fermion circuits.

As is the case with hybrid circuits without feedback (measurement outcomes are not used to inform future operations), the transition between quantum phases is only apparent in quantities nonlinear in the ensemble of quantum trajectories. In our examples, these nonlinear quantities can be spin-glass order parameters or entanglement measures. However, the spin-glass order can be converted into ferromagnetic order (a linear observable) via feedback: we show that one layer of single-qubit unitaries, conditioned on the bulk measurement outcomes, can be applied to the boundary state to ensure that the resulting density matrix averaged over trajectories has long-range order that is observable. This constitutes a nontrivial quantum channel on a 2D array producing a long-range entangled mixed state in 1D; as in [149], it relies on measurement and unitary feedback, though the resulting mixed state is likely difficult to generate using only operations on the 1D system.

This setup can be readily generalized from Pauli X, Y, and Z measurements to arbitrary single-site projective measurements. We find that such measurements in $2+0$D map in general to Gaussian fermionic hybrid circuits in $1+1$D. This mapping allows us to import the results about entanglement phases generated by such circuits (for example [101, 102, 154]) onto the measurement-induced entanglement on the boundary state of the toric code. Even in the general on-site measurement setup, the phases with area-law and logarithmic scaling of entanglement persist, albeit with distinct phase boundaries and transitions.

This chapter is organized as follows: in Section 2.2, we introduce the measurement setup for the toric code. We then map the stabilizer configurations after measurements to the completely packed loop model with crossings (CPLC) and summarize relevant results in Section 2.3. In Section 2.4, we relate specific order parameters in the loop model to entanglement induced by measurements between different regions. In Section 2.5, we explain how the mapping leads to distinct entanglement patterns in unmeasured boundary qubits and demonstrate how the setup can be mapped to a $1+1$D hybrid circuit. In Section 2.6, we show how the presence or absence of a certain spin-glass order distinguishes the two phases. Furthermore, we describe a simple adaptive protocol that modifies the boundary state and enables identification of the two phases based on linear order parameters of the state. In Section 2.7 we analyze general single-qubit measurements (beyond Pauli) on the toric code and map the resulting states to Gaussian tensor networks and Gaussian hybrid circuits. This section also contains a tensor network representation of the

toric code ground state via parton construction, which may be of independent interest. Finally, in Section 2.8 we conclude with a discussion of our results, including relations to the underlying sign structure and MBQC universality of the resource state.

2.2 Setup

2.2.1 Toric code/Plaquette model

The toric code is a lattice model of spin-1/2 degrees of freedom on the edges of a square lattice [110] which consists of commuting terms in its Hamiltonian called stabilizer operators. The toric code has two types of stabilizer operators: star (s) and plaquette (p) operators,

$$H_T = -\sum_s \prod_{j\in s} X_j - \sum_p \prod_{j\in p} Z_j \tag{2.1}$$

A closely related model is Wen's 2D plaquette model [233], where the spin-1/2 degrees of freedom are located at the vertices of a square lattice, and the Hamiltonian consists of only one type of four-body stabilizer for every star s and plaquette p, on the 45°-rotated lattice (see Fig. 2.2a):

$$H_W = -\sum_{a\in p,s} X_{a+\hat{y}} Z_{a+\hat{x}} X_{a-\hat{y}} Z_{a-\hat{x}} \tag{2.2}$$

These two models can be transformed into each other using a single layer of local Hadamard gates arranged on one (say, B) of the two sublattices (A and B) of the square lattice in the plaquette model. Sublattice A represents spins on vertical edges, while sublattice B represents spins on horizontal edges of the toric code in Fig. 2.2b. These gates interchange $X \leftrightarrow Z$ on the B sublattice, which interchange the plaquette model and toric code, as depicted in Fig. 2.2a. On a torus defined by identifying the boundaries along the x and y directions as marked in Fig. 2.2b, the toric code has four degenerate ground states, labeled by ± 1 eigenvalues of the logical operators O_1', O_2'. These logical operators can be obtained by applying the previously mentioned Hadamard gates to Wen's logical operators O_1 and O_2, which consist of strings of Pauli-Z and Pauli-X operators as depicted in Fig. 2.2b.

Any eigenstate $|G\rangle$ of the toric code admits an exact free-fermion parton construction [233, 111] defined as follows. The Hilbert space of each qubit on site i can be enlarged into that of four Majorana fermions $\gamma_{i,s} = \{\gamma_{i,1}, \gamma_{i,2}, \gamma_{i,3}, \gamma_{i,4}\}$ along the edges connected to

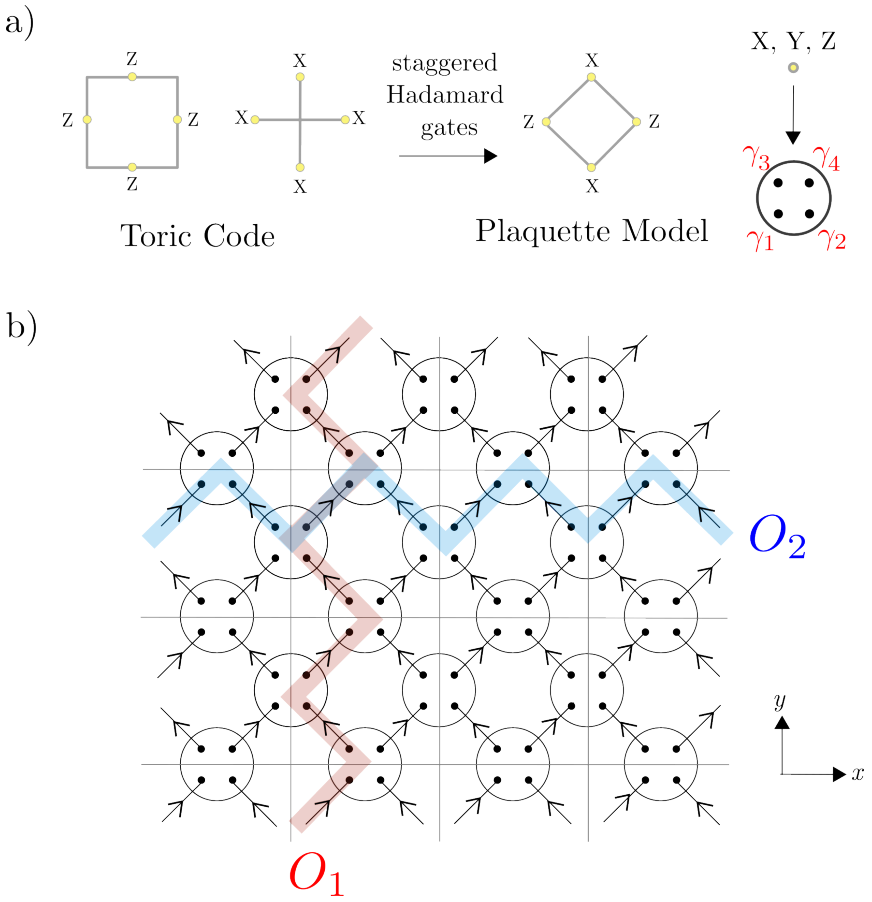


Figure 2.2: (a) Toric code stabilizers can be converted to Wen plaquette stabilizers via staggered Hadamard gates. The Wen plaquette admits a parton construction in which each qubit is split into four Majorana fermions subject to a constraint (represented by the circle in the right subfigure). (b) A ground state of the plaquette model can be constructed by projecting a free-fermion state consisting of Majorana dimers into the physical qubit Hilbert space.

i, followed by a projection onto the original qubit Hilbert space. Consider the free-fermion state $|\psi\rangle_{\text{free}}$ such that $i\gamma_{i,s}\gamma_{j,s'} = 1$ when $(i,s),(j,s')$ are on the same edge. To return to the qubit Hilbert space, we must project the Majorana state to the $+1$ sector of the operator $D_j = \gamma_{j,1}\gamma_{j,2}\gamma_{j,3}\gamma_{j,4}$:

$$|G\rangle = \prod_j \left(\frac{1+D_j}{2}\right) |\psi\rangle_{\text{free}}\,. \tag{2.3}$$

Note that the initial free-fermion state of the two Majorana modes on neighboring vertices can be oriented in two different ways, depending on whether we take the $+1$ eigenstate of $\pm i\gamma_{i,s}\gamma_{j,s'}$. Different orientations (which are marked by $s \to s'$ to indicate $+i\gamma_s\gamma_{s'}$ in Fig. 2.2b) determine the particular eigenstate up to a global phase. In particular, the ground space of the toric code on a torus is four-dimensional; one representative ground state is described in our convention by the orientation shown in Fig. 2.2b, where the same orientation is taken along all 45° lattice lines. Different logical sectors (choices of $O_1, O_2 = \pm 1$) of the plaquette-model ground space can be represented by flipping all the orientations of the links ij along the nontrivial loops. We will focus our attention on the ground state defined by the orientation shown in Fig. 2.2b, which corresponds to $O_1, O_2 = +1$.

2.2.2 Measurement setup

First we consider the case of measuring a subset of qubits M in the toric code in either the Z, Y, or X basis, with respective probabilities $(1-q)(1-p)$, p, and $q(1-p)$, which we call the (p,q) measurement protocol. Our objective is to analyze the entanglement structure and order in the remaining unmeasured qubits M^c after the measurements on M are performed. The target quantities of interest are averaged over all realizations of both measurement configurations and their outcomes. In a later section, we generalize to the case of measuring along any direction on the Bloch sphere.

Due to the equivalence between the toric code and plaquette models via a Hadamard transformation on one sublattice, the (p,q) scheme for the toric code is equivalent to the (p,q) scheme on the A sublattice and the $(p,1-q)$ scheme on the B sublattice for the plaquette model.

In the plaquette model, Pauli operators on site j correspond to Majorana-fermion

bilinear operators

$$\begin{aligned} X_j &= i\gamma_{j,1}\gamma_{j,2} = i\gamma_{j,4}\gamma_{j,3} \\ Y_j &= i\gamma_{j,2}\gamma_{j,3} = i\gamma_{j,4}\gamma_{j,1} \\ Z_j &= i\gamma_{j,1}\gamma_{j,3} = i\gamma_{j,2}\gamma_{j,4}, \end{aligned} \tag{2.4}$$

where the equalities on the right follow from the physical-Hilbert-space condition ($D_j = 1$).

2.2.3 Stabilizers and measurement

We first provide a brief overview of the Majorana stabilizer formalism specialized to our setting. The set of stabilizer generators $\mathcal{G}$ is a set of products of Majorana fermions which are independent and mutually commute with each other. This set generates the stabilizer group $\mathcal{S}$. In a Hilbert space of dimension 2^N, a set $\mathcal{G}$ with exactly N generators uniquely defines the common eigenvector $|\psi\rangle$ of any operator generated by $\mathcal{G}$, such that $s\,|\psi\rangle = |\psi\rangle$ for all $s \in \mathcal{S}$.

If we measure the state $|\psi\rangle$ with an operator P which is a product of Majorana fermions, the resulting state is still a Majorana stabilizer state and can be updated efficiently [75, 1]. There are two cases to consider. If P commutes with all the stabilizer generators $g \in \mathcal{G}$, the measurement has no effect on the state, and the measurement outcome can be inferred from the sign of the operator in $\mathcal{S}$, i.e., whether $\pm P \in \mathcal{S}$. If P anticommutes with some of the stabilizer generators, the measurement outcome is ± 1 with equal probability. We also have to modify the set of generators $\mathcal{G}$: first we select one of the anticommuting generators, denoted by g_0, and multiply g_0 with the remaining anticommuting generators. Next, we replace g_0 in $\mathcal{G}$ by either $\pm P$ depending on the measurement outcome, so that the new stabilizer set becomes $\{\pm P\} \cup \{g_0 g_i |\ \forall i \neq 0, g_i \text{ anticommutes with } P\} \cup \{g_i |\ g_i \text{ commutes with } P\}$.

The stabilizer formalism offers a way to confirm that the ground state $|G\rangle$ is the projected free-fermion state $|\psi\rangle_{\text{free}}$, stabilized by two-point Majorana-fermion operators: $i\gamma_{j,s}\gamma_{i,s'}$. The action of projection operators on the stabilizers exclusively modifies adjacent Majorana pairs by multiplying them together to form string operators. By applying all the projections, these string operators eventually become closed loops, forming Majorana loop operators which act as stabilizers for the ground state in Wen's model. Our goal is to measure Pauli operators corresponding to different two-point Majorana operators on every site. Since these are physical qubit operators, they commute with the projection operator, and hence we can first consider their effect on the free-fermion state before applying the projection operator at the end.

We graphically track the free-fermion state updates by connecting Majorana fermions with a line when they form a stabilizer operator together. When $i\gamma_j\gamma_i$ is measured, there are two possible outcomes: (a) if there is already a connection between γ_i and γ_j in the initial state, no further updates are required; and (b) if these two Majorana fermions are connected to other Majoranas (e.g., γ_i is connected to γ_k and γ_l is connected to γ_j), the update will connect γ_j to γ_i, and the other Majoranas will be connected accordingly (e.g., γ_l to γ_k), as shown in Fig. 2.3a.

The signs of stabilizers and measurement outcomes can be tracked and updated by using arrows on Majorana pairings, as illustrated in Fig. 2.3a. However, the signs will not be important when computing entanglement quantities or spin-glass order parameters in the case of X, Y, and Z measurements. In the next sections, we suppress the arrow notation for signs and return to the task of sign tracking when discussing the linear order parameter in Sec. 2.6 and on-site measurements in general directions in Sec. 2.7.

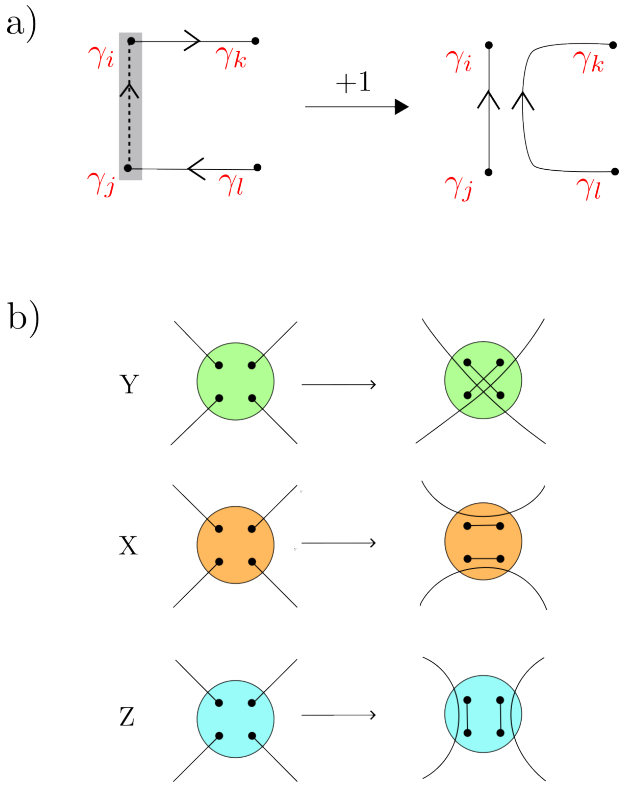


Figure 2.3: (a) The update process for measuring the Majorana bilinear $i\gamma_j\gamma_i$ for a free-fermion state which is a tensor product of Majorana dimers. The $+1$ refers to the measurement outcome, which sets the arrow direction in the final state. If the measurement outcome is -1 the final arrows need to be reversed. (b) Updated pairings from measuring X, Y, and Z Pauli operators on each site of the plaquette model. The updated pairings need to be tiled to form the global dimer state. We have neglected the sign tracking of the dimer pairs, which depends on the measurement outcomes.

2.3 Completely packed loop model

Measuring Pauli operators on each site generates three different patterns of pairings (Fig. 2.3b), and measuring all qubits tiles these patterns and results in different configurations of loops on a square lattice. On the two different sublattices of the square lattice, the factors q and $1-q$ must be swapped, to reflect the staggered measurement scheme of the plaquette model.

Consider a configuration of measurements or tilings, with N_x, N_y, and N_z the numbers of X, Y, and Z measurements performed. Such a configuration has probability $W_{\mathcal{C}} = p^{N_y}[(1-p)q]^{N_x}[(1-p)(1-q)]^{N_z}$, which leads to the partition function $Z = \sum_{\mathcal{C}} W_{\mathcal{C}}$. The model and partition function are known as the completely packed loop model with crossings (CPLC), whose properties have been extensively studied in [160]. We now review its important properties relevant to the questions addressed in this chapter.

In [160] the authors found that the phase diagram consists of a short-loop phase and a long-loop "Goldstone" phase, which are separated by a phase transition (see Fig. 2.4). This model can be described by the replica limit $n \to 1$ of a $\mathbb{Z}_2$ lattice gauge theory coupled to $O(n)$ matter field. Its continuum description is a sigma model, which is massive in the short-loop regime and massless in the Goldstone phase [160].

We focus on two order parameters which distinguish the phases. First, we consider the *watermelon correlation functions* $G_k(i,j)$, which denote the probability that k distinct strands connect points i and j, where k is even for the CPLC model. For instance, G_4 is the probability that two nodes are connected by four distinct strands. Using renormalization-group (RG) techniques on the sigma model, [160] found that in the Goldstone phase

$$G_k(i,j) \sim \frac{C_0}{\ln\left(d_{ij}/r_0\right)^{k(k-1)}} \tag{2.5}$$

where d_{ij} is the distance between i and j and C_0, r_0 are non-universal constants. In the short-loop phases, on the other hand, the watermelon correlators decay as $G_k(i,j) \sim e^{-d_{ij}/\xi}$, with correlation length ξ.

Next, we consider the *spanning number* defined for a CPLC model on a cylinder with two circular open boundaries. The spanning number counts the number of strands that connect the upper and lower boundaries. Reference [160] found that in the Goldstone phase, the average spanning number scales with system size L as

$$n_s \approx \frac{1}{2\pi}\left(\ln\frac{L}{L_0} + \ln\ln\frac{L}{L_0}\right), \tag{2.6}$$

whereas it asymptotes to 0 in the short-loop phase.

To explore the entanglement properties of the toric code after measurements and their connection to the phase transitions in the loop model, we need to leave some qubits unmeasured, as measuring all qubits results in a trivial pure product state. Three scenarios are considered (see the schematic description in Fig. 2.4):

(I) *Measuring all but two qubits in the bulk.* In Sec. 2.4.1 we show that the entanglement induced between the two unmeasured qubits is directly related to the watermelon correlation function.

(II) *Measuring all but two boundaries.* In Sec. 2.4.2, we observe that the induced entanglement between the two boundaries of the cylinder is directly related to the "spanning number" order parameter discussed in this section. Accordingly, in the short-loop phase, the entanglement is asymptotically zero, while in the Goldstone phase, it exhibits logarithmic scaling with the system size.

(III) *Measuring all but a single boundary.* We show in Sec. 2.5 that in this case the entanglement between contiguous bipartitions of the unmeasured boundary exhibits a phase transition between area-law and logarithmic scaling, reflecting the underlying loop-model configurations.

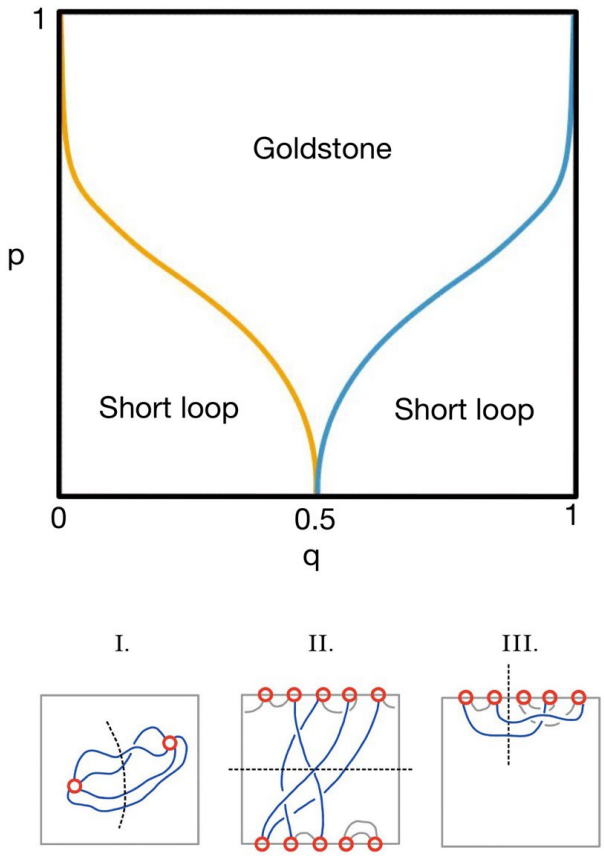


Figure 2.4: The schematic phase diagram of the completely packed loop model with crossings (CPLC) [160]. The diagram distinguishes between different phases based on three order parameters, highlighted in the panel below. These order parameters correspond to loop configurations that connect points marked by red dots. The first order parameter (I) quantifies the probability of four distinct strands connecting two red points on a torus, referred to as the "watermelon correlation function." The second order parameter (II) measures the expected number of strands connecting the top and bottom boundaries on a cylinder, known as the "spanning number." The third order parameter (III) measures the expected number of strands connecting two partitions of the top boundary on a cylinder with a fixed boundary condition at the bottom. All three quantities govern measurement-induced entanglement in the toric-code state.

2.4 Measurement-induced entanglement between two distant regions

In this section, we establish a connection between the average measurement-induced entanglement (MIE) of two distant unmeasured regions and various order parameters within the CPLC model. The MIE has been related to the underlying sign structure of the measured wavefunction in [145], and we comment more on this in the concluding discussion.

2.4.1 MIE between two unmeasured qubits

We first demonstrate that the measurement-induced entanglement (MIE) between two unmeasured qubits at sites i and j is equivalent to the watermelon correlation $G_4(i,j)$.

Recall that measuring a qubit specifies a given pairing for the four Majoranas associated with the qubit. After all pairings at all sites except for i,j are specified, we must implement the projection operators D_k on every site, as in Eq. 2.3. Crucially, any *closed-loop Majorana stabilizer commutes with the projection operator* and is thus shared by both the free-fermion and the projected state in Eq. 2.3. However, if two qubits are left unmeasured, then some Majorana stabilizers may be open strands ending at the unmeasured sites. In this case the projection operator has a significant effect on the final stabilizers and hence the entanglement between the unmeasured qubits.

To compute the measurement-induced entanglement, we analyze the three ways (Fig. 2.5) in which Majorana stabilizer strands terminate at the two vertices i,j. (Any closed loop not coincident with i and j does not contribute any entanglement.) Denote a stabilizer strand connecting Majoranas $\gamma_{i,s}$ and $\gamma_{i',s'}$ as $(i_s i'_{s'})$. We suppress the sign information of the stabilizer in this notation. The three classes of configurations are

(a) Each strand ends on Majoranas on the same vertex, i.e. we have 2 $(i_s i_{s'})$ and 2 $(j_s, j_{s'})$ pairings.

(b) Two strands end on the same vertex and two strands end on different vertices, i.e. there are 1 $(i_s i_{s'})$, 2 $(i_s, j_{s'})$, and 1 $(j_s, j_{s'})$ pairings.

(c) All four strands terminate on different vertices, i.e. there are 4 $(i_s, j_{s'})$ pairings.

Once we impose the local projection operators $D_i = (i_1 i_2 i_3 i_4)$ and $D_j = (j_1 j_2 j_3 j_4)$, the stabilizer generators need to be updated. Furthermore, to compute the entanglement between the two unmeasured qubits, we choose a canonical gauge for the stabilizers [159, 141]; for a given bipartition, stabilizer generators restricted to either subsystem are independent.

In this canonical form, the entanglement across a bipartition is proportional to the number of stabilizers shared between both parties.

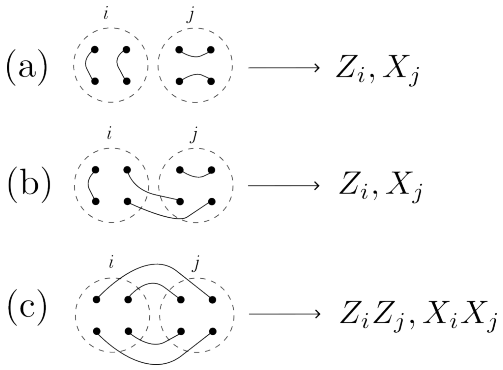


Figure 2.5: Three possible pairings of unmeasured qubits i, j after all other qubits are measured. Left: stabilizer strands prior to physical-qubit Hilbert-space projections. Right: stabilizers after projections, in canonical form. In writing the stabilizers as Pauli strings, we have assumed that the sites i, j are in the same sublattice. Configurations (a, b) are exclusively supported on i or j and do not contribute entanglement, while configuration c contributes one bit of entanglement.

We show examples of the stabilizer update into its canonical form in Fig. 2.5. In the right column in Fig. 2.5, we show the stabilizer generators obtained from these patterns of strands, set in their canonical form such that the stabilizer generators restricted to either i or j are independent. As can be seen, the number of such independent *connecting* stabilizers is $0, 0$, and 2, respectively, in the three types of Majorana pairings a, b, and c. Thus, only configuration c contributes one bit of entanglement, and the average measurement-induced entanglement (MIE) generated between any two unmeasured qubits on i and j is exactly given by the probability of four distinct loops connecting i and j in the CPLC model, i.e. the watermelon correlation function defined in the previous section,

$$\langle S_{\mathrm{MIE}}(i,j)\rangle = (G_4^{\mathrm{CPLC}}(i,j))\ln 2. \tag{2.7}$$

Hence it follows from the results of [160], as quoted in Eq. 2.5, that the averaged MIE is long-ranged in the Goldstone phase and short-ranged in the short-loop phase.

2.4.2 MIE between two unmeasured boundaries

Now we consider the toric code on a cylinder and explore the effects of bulk measurements on the boundary chains of qubits; in particular, we focus on the scenario where

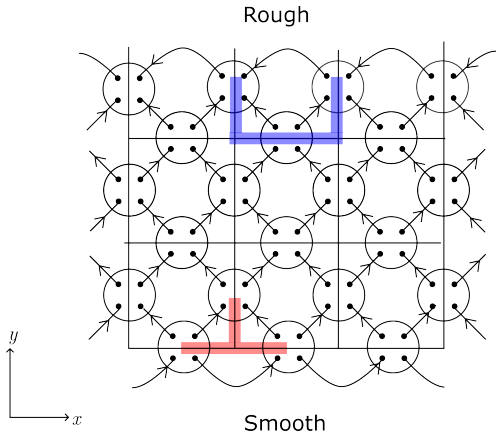


Figure 2.6: Boundary conditions of the parton ground state of the toric code on a cylinder. The rough and smooth boundary conditions on two open edges of the cylinder are shown, along with the respective truncated stabilizers on the boundaries. The x direction is taken to have periodic boundary conditions.

both circular boundaries of the toric code are left unmeasured. For the purposes of this section, the exact boundary conditions do not matter, so we defer a discussion of the exact boundary conditions and the exact mapping of this scenario to the loop model to the next section. Here we simply quote the final result: under the Majorana-loop mapping, the average entanglement between these two boundaries can be directly mapped to the "spanning number" in the loop model, as illustrated in Fig. 2.4II.

This can already be motivated from the discussion in the earlier subsection, where we showed that the entanglement between two remote regions of the toric code corresponds to open strands connecting the regions in the CPLC model. However, we must also transform the stabilizers to their canonical forms in order to directly count their entanglement contribution. In Appendix A.1.2 we show that in this geometry, if there are $n \geq 2$ such strands connecting the top and bottom boundaries in the loop model, we get $n-2$ independent Majorana stabilizer generators in their canonical form connecting the top and bottom boundaries. The average n is just the spanning number of the loop model, so we get the following correspondence,

$$\langle S_{\mathrm{MIE}} \rangle = \frac{(n_s - 2)\ln 2}{2}. \tag{2.8}$$

This correspondence holds true only for $n_s \geq 2$; otherwise, we have $S_{\mathrm{MIE}} = 0$. As noted in Eq. 2.6, the average spanning number n_s scales logarithmically with L in the Goldstone phase and asymptotes to zero in the short-loop phase.

2.5 Measurement-induced phase transition in the boundary

In this section, we investigate the occurrence of measurement-induced phase transitions in a 1D chain of qubits on the boundary of a toric-code state, where the remaining qubits are measured in random local Pauli bases (as depicted in Fig. 2.1). Many of the findings in this section can be generalized to different topologies (such as torus, cylinder, or plane) and partitioning schemes of the unmeasured 1D chain.

We begin by examining a toric code implemented on a cylinder with two open circular boundaries. The boundary stabilizers are truncated, resulting in two types of boundary conditions: "rough" or "smooth." The rough condition arises when the plaquette stabilizers are truncated, while the smooth condition occurs when the star stabilizers are truncated. In Fig. 2.6, the truncated toric code represented in the parton picture exhibits a simple form, where no distinction between rough and smooth is evident. Consequently, it is unnecessary to specify the type of boundary condition for the entanglement analysis. Moreover, various boundary conditions for the surface code can be created by measuring the toric-code state on a torus, followed by applying a single layer of local unitary updates based on the measurement outcomes. For instance, to achieve a surface code with smooth boundary conditions, one should perform Pauli-X measurements along a horizontal line on sublattice A. Conversely, a surface code with rough boundary conditions can be formed by conducting Pauli-Z measurements along a horizontal line on sublattice B. Applying both of these measurement patterns results in the surface code illustrated in Fig. 2.6, featuring both rough and smooth boundary conditions. For concreteness, we will hereafter assume the pair of rough and smooth boundary conditions for the surface code.

In this section, we show that the boundary state after bulk measurements is precisely the state generated by a hybrid circuit in $1+1$D composed of measurements and unitaries. While this identification is general, for the particular case of the toric code, the corresponding hybrid circuit is a $1+1$D Ising-symmetric circuit [195, 192]. Our mapping is motivated by MBQC, whereby a circuit can be effectively realized by single-site measurements on a resource state. However, the toric code is not a universal resource state [33]; thus, single-site measurement on the toric code cannot represent all circuits.

2.5.1 Measured toric code as a $1+1$D hybrid circuit

Starting from the parton representation of the toric code, measuring each qubit modifies the pairings of four neighboring Majoranas, as described in Sec. 2.2. These pairings can be

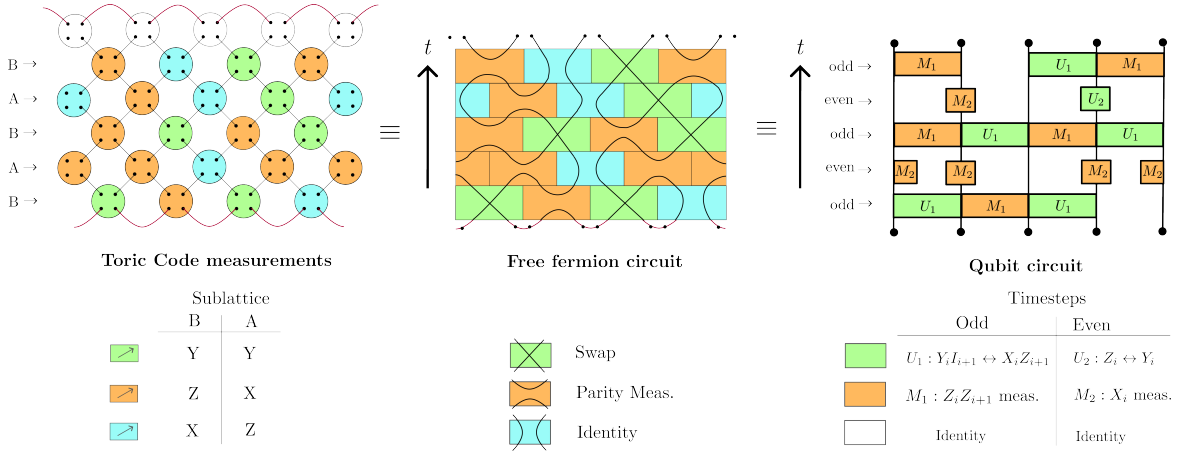


Figure 2.7: Pauli-$X/Y/Z$ measurements on the bulk, with the boundary chain of qubits left unmeasured (left panel), correspond to a hybrid circuit with measurement and unitary gates. The measurement model is equivalent to a staggered measurement protocol on the Wen plaquette model. The plaquette model generates Majorana pairing patterns that can be interpreted as world lines of Majorana fermions undergoing a free-fermion hybrid circuit (middle panel). By Jordan-Wigner transformation, the same circuit can be identified with an Ising-symmetric measurement and unitary circuit on qubits (right panel).

interpreted as world lines of two Majoranas undergoing circuit operations. In particular, the corresponding circuit consists of the following gates acting on two neighboring Majorana fermions: swap gate, identity gate, and fermion parity measurement, as illustrated at the bottom of Fig. 2.7. Thus, starting from an $N \times N$ toric-code state on a cylinder, the bulk measurements realize a depth-N free-fermion circuit on $2N$ Majorana fermions (Fig. 2.7, middle). Note that for a fixed configuration of measurement bases, different measurement outcomes correspond to hybrid circuits in $1+1$D differing only in the signs of the Majorana stabilizers. However, as noted earlier, the observables we are interested in—entanglement and spin-glass order—are not sensitive to these signs and only depend on the worldline connectivity. Thus, for these observables, the hybrid circuits for a given measurement-basis configuration are equivalent, regardless of the outcomes.

After all the projections onto the physical qubit Hilbert space are imposed, the Majorana hybrid circuit maps via Jordan-Wigner transformation to a depth-N hybrid circuit of local unitaries and local measurements on a one-dimensional system of N qubits. This mapping is shown explicitly in Fig. 2.7. We note that this mapping is somewhat subtle, as the parton construction and Jordan-Wigner transformation are two different mappings from one qubit to respectively four and two Majorana fermions. Briefly, the reason why the

claimed mapping works is because at the top boundary, the top two of the four Majoranas per parton decomposition are always paired in a nearest-neighbor dimer state (Fig. 2.7, left), so the physical qubit state is solely determined by the bottom two Majoranas per site via the standard Jordan-Wigner mapping (see Appendix A.2 for details).

The Majorana-fermion parity measurements realize the measurements of either the neighboring $Z_k Z_{k+1}$ or on-site X_k measurement in the qubit circuit, depending on which sublattice the measurements are performed:

$$\begin{aligned} M_1 = i\tilde{\gamma}_{k,2}\tilde{\gamma}_{k+1,1} &= Z_k Z_{k+1} \\ M_2 = i\tilde{\gamma}_{k,1}\tilde{\gamma}_{k,2} &= X_k. \end{aligned} \tag{2.9}$$

Similarly, the Majorana swap gate implements either a two-qubit unitary U_1 or on-site unitary U_2:

$$\begin{aligned} U_1 &: Y_k I_{k+1} \leftrightarrow X_k Z_{k+1} \\ U_2 &: Z_k \leftrightarrow Y_k. \end{aligned} \tag{2.10}$$

These circuit operations preserve an Ising $\mathbb{Z}_2$ symmetry $\prod_i X_i$. The origin of this Ising symmetry of the hybrid circuit is the fact that the $\prod_i X_i$ string operator, supported on the rough boundary, takes a definite value for the initial surface-code state, and any bulk measurement away from the boundary commutes with this boundary operator.

The qubit circuit corresponding to the (p, q) measurement protocol on the toric code (TC) state can be explicitly defined as follows (see also Table 2.1):

(1) In odd times, perform one of the following operations on each pair of neighboring qubits: the two-qubit unitary U_1 (TC measurement along Y) with probability p, the identity operation with probability $(1-p)q$ (TC measurement along X), and the M_1 measurement with probability $(1-p)(1-q)$ (TC measurement along Z).

(2) In even times, perform one of the following operations on each qubit: the on-site unitary U_2 with probability p (TC measurement along Y), the on-site measurement M_2 with probability $(1-p)q$ (TC measurement along X), and the identity operation with probability $(1-p)(1-q)$ (TC measurement along Z).

2.5.2 Bipartite entanglement in the unmeasured boundary

By identifying the bulk Majorana pairings with the classical loop model as in Sec. 2.3, we can establish a direct mapping between the entanglement across a bipartition of the

Sublattice (timesteps)	Measurement TC (WP)	Probability	Qubit Circuit
B (odd)	Y (Y)	p	U_1
	X (Z)	$(1-p)q$	Identity
	Z (X)	$(1-p)(1-q)$	M_1
A (even)	Y (Y)	p	U_2
	X (X)	$(1-p)q$	M_2
	Z (Z)	$(1-p)(1-q)$	Identity

Table 2.1: Circuit mapping of the (p, q) measurement model on the toric code (TC). For completeness, the corresponding staggered measurement protocol for the Wen plaquette (WP) model is mentioned in parentheses.

boundary state and a specific quantity depicted schematically in Fig. 2.4III within the loop model. This quantity, which counts the number of strands connecting two parts of the boundary chain, diagnoses the long-range correlations in the Goldstone phase of the loop model.

We also show in Appendix A.1.1 that for the one-boundary setup, there is a one-to-one correspondence between the Majorana strands and the canonical stabilizer generators after implementing projections. Specifically, a configuration with n open strands connecting the two parts of the boundary corresponds to a quantum state with n canonical stabilizer generators connecting the two parts, thereby contributing $\frac{n \ln 2}{2}$ units of entanglement. The bipartite entanglement between two parts of the boundary thus acts as an order parameter for the phase transition in the loop configurations of the CPLC model.

In the Goldstone phase of the CPLC, [195] found that the entanglement between a contiguous subregion A of the qubit chain has a logarithmic scaling with a correction. Therefore, the entanglement of a contiguous subregion A of the unmeasured boundary chain of the toric code also satisfies the same entanglement scaling in the Goldstone phase,

$$\langle S_A \rangle \approx \frac{\ln(2)}{2} \left(\# \ \ln |A| + \frac{1}{4\pi} (\ln |A|)^2 \right), \tag{2.11}$$

while it obeys an area law in the short-loop phase.

2.6 Long-range order in the boundary state

2.6.1 Spin-glass order parameter

As we showed in the previous section, the bulk measurements performed on the toric code can be mapped to the Ising-symmetric hybrid circuits studied by Sang et al. [195, 192]. Given only ZZ measurements, the steady state is a "random GHZ" state characterized by a random spin configuration superposed with the flipped configuration. This is also known as a spin-glass state, and the spin-glass order is captured by the Edwards-Anderson order parameter:

$$O = \frac{1}{L}\sum_{i,j}^{L} \langle\psi|Z_i Z_j|\psi\rangle^2 \tag{2.12}$$

For spin-glass order, $O \sim L$, whereas for paramagnetic order, $O \sim O(1)$.

Reference [192] studied the phase diagram of hybrid circuits involving ZZ and X measurements and random Ising-symmetric Clifford unitaries, and a stable spin-glass phase was found. The toric-code measurements map to a subset of symmetric Clifford unitaries—namely the free-fermion operations defined above—and for this restricted class we provide a fermionic perspective on the spin-glass order parameter and the extent of the spin-glass phase.

The central object in the spin-glass order parameter is $Z_i Z_j$, which for a stabilizer state can take three values (± 1 or 0). It maps via Jordan-Wigner transformation to

$$Z_i Z_j = i\gamma_{i,2}\left(\prod_{k=i+1}^{k=j-1} i\gamma_{k,1}\gamma_{k,2}\right)\gamma_{j,1}. \tag{2.13}$$

This string of Majoranas is nonzero if and only if all Majoranas within the interval (i, j) are paired amongst themselves. (If any Majorana within the interval is paired with one outside, that dimer will anticommute with the above string and render its expectation value zero.)

In the short-loop phase, configurations are composed of loops with a characteristic size ξ, which is independent of L. Hence, the probability that all Majoranas within an interval (i, j) are paired up is independent of $|i - j|$ for $|i - j| \gg \xi$. For $q < 1/2$, this probability is a nonzero $O(1)$ number independent of $|i - j|$, and thus the spin-glass order parameter is extensive in the $q < 1/2$ short-loop phase.

2.6.2 Linear order parameter from adaptive circuits

Due to the equal probabilities of Z_iZ_j having opposite signs ± 1, its average value is zero, making it necessary to use nonlinear order parameters such as the Edwards-Anderson order parameter defined in Eq. 2.12. However, in experimental setups, measuring nonlinear order parameters is generally challenging and requires post-selection of the measurement outcomes. In practice it is more feasible to access expectation values of operators like $\mathrm{Tr}(\rho O)$ which are linear in the density matrix ρ, the ensemble of all measurement trajectories.

Here we detail an efficient protocol for converting the (nonlinear) spin-glass order into a linear order parameter. Our protocol employs an adaptive strategy that applies a layer of local unitaries on the boundary conditioned on the bulk measurement outcomes. The main objective of this adaptive unitary layer is to transform all the $\pm$ values of Z_iZ_j stabilizer generators into positive values, thus eliminating the issue of sign cancellation and converting the spin-glass order into long-range ferromagnetic order, which is linear in the state and can be accessed experimentally. The protocol consists of two main parts: (1) identify Z_iZ_j stabilizers and determine their signs, and (2) obtain single-site unitaries that can correct the negative signs to positive signs.

Note that one approach for identifying and correcting the sign of Z_iZ_j operators in the stabilizer group is to simulate the entire evolution classically using the stabilizer formalism, with the knowledge of all the $O(L^2)$ bulk measurement outcomes. However, we propose a simpler algorithm that only requires access to the directions and outcomes of measurements within a correlation length $O(\xi)$ from the boundary, i.e. $O(L\xi)$ measurements.

The algorithm is as follows:

A. Z_iZ_j generator graph construction: We construct a graph whose vertices are the boundary qubit sites i and which has an edge ij if $\pm Z_iZ_j$ is in the stabilizer group. By knowing the positions of X, Y, and Z measurements, we can obtain the corresponding configuration of Majorana strands. If for any $i < j$, $\gamma_{i,2}$ is paired with $\gamma_{j,1}$, we draw an edge ij in the graph. As per Eq. 2.13, this pairing implies a Z_iZ_j stabilizer if and only if all the Majorana fermions in the interval (i, j) are paired internally. This can be checked for all the Majorana fermions in (i, j) from the loop configuration. If there are strands that exit the interval (i, j), then we erase the edge ij in the graph as this does not correspond to a Z_iZ_j stabilizer. For the special case $p = 0$ without loop crossings, this second step is not necessary, as all the Majorana strands must be nested in this case. Note that the graph is a tree as the Majorana strands are all independent generators.

We then obtain the sign of the Z_iZ_j stabilizers by tracking the sign of the arrows along

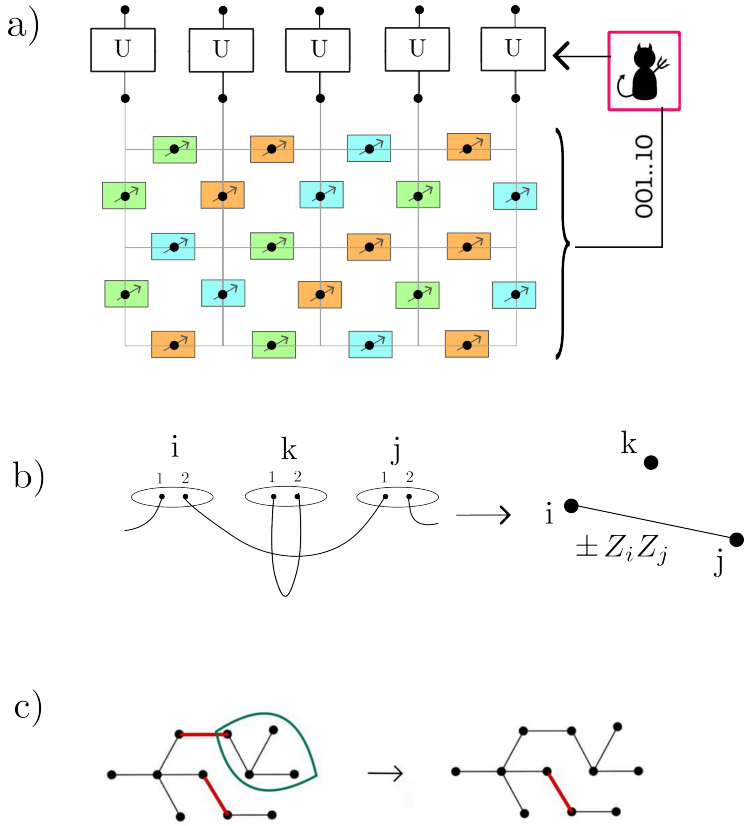


Figure 2.8: (a) To convert any (nonlinear) spin-glass long-range order of the boundary state into (linear) ferromagnetic long-range order, a layer of on-site unitaries conditioned on bulk measurement outcomes can be applied to the boundary. (b) The classical processing for the adaptive protocol involves using the distribution of Majorana strands to construct a graph in which an edge between nodes (qubits) i, j represents the existence of a $Z_i Z_j$ stabilizer. The edges are colored depending on whether the sign of the stabilizer is ± 1, which can be computed from the measurement outcomes along the strands. (c) The negative signs in the tree graph can be flipped by applying X_k on all nodes on one side of the negative edge (e.g. those encircled).

the strand $\gamma_{i,2}\gamma_{j,1}$ and all the intermediate strands $\gamma_{k,1}\gamma_{k',2}$ for $k, k' \in (i, j)$. This requires us to keep track of only $O(L\xi)$ measurement outcomes, where ξ is the correlation length corresponding to the loop size in the underlying CPLC model. If $Z_iZ_j = -1$, we color the edge ij.

B. Correcting the sign: From the previous step, we have a colored graph. Flipping a spin k (acting with the unitary X_k) changes the signs of all adjacent edges to the node k. We can correct any negative sign on an edge by flipping all nodes on one side of the edge (see Fig. 2.8c). Repeating this process for each edge allows us to flip the sign of every edge individually. To correct all the edges with a minimal number of operations, a search algorithm within the tree can be performed, which has polynomial complexity with respect to the system size.

Given access to the measurement protocol and outcomes, this algorithm (with complexity polynomial in system size) can be executed by a classical computer to determine the necessary adaptive unitary protocol, resulting in all trajectories having non-negative Z_iZ_j correlations. The resulting ensemble of trajectories ρ has long-range ferromagnetic order $\langle Z_iZ_j\rangle_\rho$ in the original spin-glass phase.

2.7 General on-site measurements

We now consider the effect of measurements along an arbitrary direction $n_xX + n_yY + n_zZ$ on the toric code. These map to statistical models that go beyond the scope of the previously discussed CPLC model (Section 2.3). In the following, we demonstrate the representation of the measured toric code using Gaussian tensor networks (GTN) and the realization of Gaussian hybrid circuits within the virtual space. Furthermore, we establish a connection with well-studied Gaussian hybrid circuits to showcase the robustness of the boundary MIE phase diagram obtained from the CPLC model in the general measurement case.

2.7.1 Tensor network representation of parton construction

The parton representation of the toric-code state can be reinterpreted as a two-dimensional tensor network composed of local tensors $|T\rangle$. These tensors consist of a single physical leg representing a qubit and four virtual legs representing the parton Majorana fermions. To construct a tensor network, contractions between spins or contractions between Majorana fermions are allowed. A contraction between spins involves projecting the two legs onto a

maximally entangled state, and a contraction between Majorana fermions sharing an edge involves projecting the two Majorana fermions onto the eigenstate of the $i\gamma_i\gamma_j$ operator with eigenvalue $+1$. The latter assumes an orientation for each virtual bond which must be specified.

We now describe the tensor-network representation of the projected parton states of the toric code. First, we introduce the projection tensor that maps four Majorana fermions to one spin:

$$P = \frac{|G_1\rangle |\uparrow\rangle + |G_2\rangle |\downarrow\rangle}{\sqrt{2}}. \tag{2.14}$$

Here $|G_1\rangle$ and $|G_2\rangle$ are two fermionic stabilizer states of the four Majoranas, where $|G_1\rangle$ is stabilized by $i\gamma_1\gamma_3$ and $i\gamma_2\gamma_4$, and $|G_2\rangle = i\gamma_1\gamma_2 |G_1\rangle$. The orientation for bonds between tensors is specified in Fig. 2.9(a), and corresponds to the orientation of Majorana pairs described in Section 2.2.

2.7.2 Measured toric code as a Gaussian tensor network

Measuring a qubit along the $\vec{n}$ axis and obtaining outcome ± 1 corresponds to contracting the physical leg of a tensor with the qubit state $|\psi_{\pm\vec{n}}\rangle$, the eigenstate of the operator $\vec{n}\cdot\vec{\sigma} = n_x X + n_y Y + n_z Z$ with eigenvalue ± 1. After contraction, the resulting projection tensor takes the form of a Gaussian fermionic tensor on the remaining Majorana legs:

$$\begin{aligned} |F_{\vec{n}}\rangle &= \frac{|G_1\rangle \langle\psi_{\vec{n}}| |\uparrow\rangle + |G_2\rangle \langle\psi_{\vec{n}}| |\downarrow\rangle}{\sqrt{2}} \\ &= \frac{e^{i\phi}\cos\frac{\theta}{2} |G_1\rangle + \sin\frac{\theta}{2} |G_2\rangle}{\sqrt{2}} \end{aligned} \tag{2.15}$$

Here, θ and ϕ represent the spherical coordinates of the unit vector $\vec{n} = (n_x, n_y, n_z) = (\sin\theta\cos\phi, \sin\theta\sin\phi, \cos\theta)$. The fermionic operators F^1 and F^2 which stabilize $|F_{\vec{n}}\rangle$ (meaning that $F^{1,2} |F_{\vec{n}}\rangle = |F_{\vec{n}}\rangle$) are:

$$\begin{aligned} F^1 &= i\gamma_1\gamma_3\cos\theta + i\gamma_1\gamma_2\sin\theta\cos\phi + i\gamma_1\gamma_4\sin\theta\sin\phi \\ F^2 &= i\gamma_2\gamma_4\cos\theta + i\gamma_4\gamma_3\sin\theta\cos\phi + i\gamma_3\gamma_2\sin\theta\sin\phi. \end{aligned}$$

The independence, commutation, and unit-square properties of the fermionic operators F^1 and F^2 can be straightforwardly demonstrated. These two fermionic operators uniquely define $|F_{\vec{n}}\rangle$ as a Gaussian state/tensor which, in fact, is the most general form for four

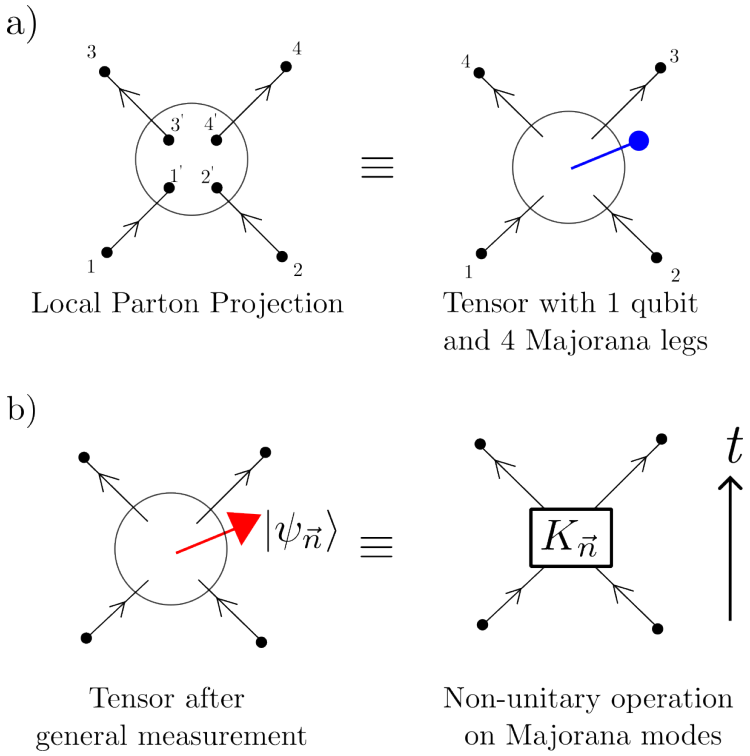


Figure 2.9: (a) The parton representation of the toric code state can be interpreted as a local tensor with one "physical" qubit leg and four "virtual" Majorana legs, along with an orientation for contraction of the virtual indices. In this case, the orientation corresponding to a toric-code ground state is specified. (b) Measuring the parton state, or equivalently, contracting the physical index with a state vector in a general direction, leads to a Gaussian state on the four "virtual" Majorana degrees of freedom. This Gaussian state can be mapped to a Gaussian operation on two Majorana fermions, and can be directly identified with a non-unitary operator.

Majorana fermions. The covariance matrix of the state (which is defined as $\Gamma_{ij} = \langle \frac{i}{2}[\gamma_i, \gamma_j] \rangle$) in the $(\gamma_1, \gamma_2, \gamma_3, \gamma_4)$ basis is:

$$\Gamma = \frac{1}{2} \begin{bmatrix} 0 & \sin\theta\cos\phi & \cos\theta & \sin\theta\sin\phi \\ -\sin\theta\cos\phi & 0 & -\sin\theta\sin\phi & \cos\theta \\ -\cos\theta & \sin\theta\sin\phi & 0 & -\sin\theta\cos\phi \\ -\sin\theta\sin\phi & -\cos\theta & \sin\theta\cos\phi & 0 \end{bmatrix}$$

In summary, each qubit measurement in the $\vec{n}$ direction results in a Gaussian tensor supported on virtual legs. In the setup where all but a top boundary of qubits are measured, all degrees of freedom except the top boundary of qubits are contracted (Fig. 2.10). We show in Appendix A.3 that the Jordan-Wigner transformation of this boundary qubit state precisely yields the Gaussian tensor-network state in which the boundary qubit projections and top row of Majorana contractions are removed (Fig. 2.10, right). Consequently, performing measurements in the bulk of the toric code and employing a Jordan-Wigner transformation at the boundary results in a remaining state described by a Gaussian tensor network.

In the measured toric code, each quantum trajectory can be represented by $\vec{n}_i$, where $\vec{n}_i$ indicates both the type and the outcome of the measurement when measuring the qubit at site i. In this notation, the type of measurement is encoded in the axis of $\vec{n}$, denoted as $e_{\vec{n}}$, while the measurement result is encoded in the direction of $\vec{n}$ along the axis. The probability of each trajectory arises from two sources. First, there is the classical probability set by the protocol, denoted as $w(e_{\vec{n}})$, representing the probability of selecting measurement axis $e_{\vec{n}}$. If the measurement bases are chosen independently from site to site, the probability of a given set of measurement bases is $w(\{e_{\vec{n}_i}\}) = \prod_i w(e_{\vec{n}_i})$. Second, the Born rule for the measurement outcome determines the probability $P(\vec{n})$ of the sign associated with $\vec{n}$ based on the measurement. The Born probability $P(\{\vec{n}_i\})$ of each trajectory $\{\vec{n}_i\}$ is the norm of the wave function represented by the tensor network. The combined probability for a particular quantum trajectory is thus $w(\{e_{\vec{n}_i}\})P(\{\vec{n}_i\})$.

This notation can be used to represent the (p, q) measurement protocol for Z, Y, or X measurements discussed in previous sections. Due to the axis-based definition of w and the fact that $e_{\vec{n}} = e_{-\vec{n}}$, the probability distribution trivially satisfies $w(e_{\vec{n}}) = w(e_{-\vec{n}})$. Therefore, the weights associated with these measurements are $w(e_Z) = w(e_{-Z}) = (1-q)(1-p)$, $w(e_Y) = w(e_{-Y}) = p$, and $w(e_X) = w(e_{-X}) = q(1-p)$. Based on the stabilizer formalism, the Born probability for each plus and minus sign is equal to $P(\vec{\sigma}) = \frac{1}{2}$. Consequently, $\sum_{\vec{\sigma}} w(e_{\vec{\sigma}}) P(\vec{\sigma}) = 1$. Also note that the classical probability satisfies:

$$\sum_{e_{\vec{\sigma}}} w(e_{\vec{\sigma}}) = w(e_X) + w(e_Y) + w(e_Z) = 1.$$

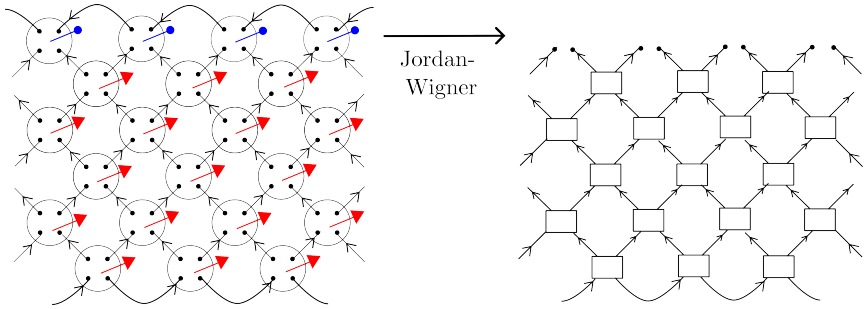


Figure 2.10: Bulk measurements of general on-site operators on the toric-code ground state can be mapped to a Gaussian circuit on a chain of the virtual Majorana degrees of freedom. The qubit state supported on the top boundary on the left-hand side maps under Jordan-Wigner transformation to the Gaussian state supported on the top boundary on the right-hand side (see Appendix A.3 for derivation).

2.7.3 Measured toric code as a Gaussian hybrid circuit

The Gaussian tensors in the virtual space can be understood as Gaussian operations that perform non-unitary transformations between the lower and upper legs of the virtual space, as illustrated in Fig. 2.9(b). Building upon this observation, we establish a correspondence between a specific class of non-unitary Gaussian circuits in 1+1D and the random Gaussian tensor networks in 2D discussed in the previous section.

These Gaussian circuits are characterized by a set of operators $\{K_{\vec{n}}\}$ and a probability distribution $\tilde{w}(\vec{n})$ associated with the operations. The operators $K_{\vec{n}}$ act on the neighboring Majorana modes γ_i and γ_{i+1} within the virtual space, resulting in non-unitary Gaussian circuits [101, 102]. More explicitly, these operators can be expressed as:

$$K_{\vec{n}} = \left(1 - n_x^2\right)^{1/4} e^{-i\alpha(\vec{n})\gamma_i\gamma_{i+1}} \tag{2.16}$$

where $\alpha(\vec{n})$ is defined by

$$e^{-2\mathrm{Re}[\alpha(\vec{n})]} = \left(\frac{1+n_x}{1-n_x}\right)^{1/2}; e^{i2\mathrm{Im}[\alpha(\vec{n})]} = \frac{in_y + n_z}{\left(n_y^2 + n_z^2\right)^{1/2}} \tag{2.17}$$

It is worth noting that when n_x is zero, the operations are unitary. However, when n_x is ± 1, the operations correspond to fermion-parity measurements. For general $\vec{n}$ with $n_x \neq 0$, the non-unitary operation $K_{\vec{n}}$ is a weak measurement of the fermion parity. The

spacetime geometry of the monitored Gaussian circuit acting on a Majorana chain is depicted on the right-hand side of Fig. 2.10, and each operation is randomly chosen from the set of operations $K_{\vec{n}}$. The probability of each trajectory is determined by two sources: the classical distribution $\tilde{w}(\vec{n}_i)$ and the Born probability $\langle\psi| K^\dagger_{\vec{n}_i} K_{\vec{n}_i} |\psi\rangle$, where $|\psi\rangle$ is the normalized wave function of the 1D fermion chain before the action of the operator $K_{\vec{n}_i}$, and the updated wave function after the action is given by $\frac{K_{\vec{n}_i}|\psi\rangle}{||K_{\vec{n}_i}|\psi\rangle||}$. Note that the classical probability should satisfy $\tilde{w}((n_x, n_y, n_z)) = \tilde{w}((-n_x, n_y, n_z))$ so as not to force any bias on the measurement outcomes.

The first step in establishing the correspondence is to relate the operator representation of the tensor $|F_{\vec{n}}\rangle$ in the tensor network to the set of operators in the Gaussian circuit. By performing explicit contractions of the tensor network, it can be shown that the operator representation of the tensor $|F_{\vec{n}}\rangle$ is $\frac{K_{\vec{n}}}{\sqrt{2}}$. This implies that the final state of the hybrid circuit and the tensor network, given the same set of directions, are equal up to a normalization factor. However, this normalization factor is only relevant for the probability of the trajectory. This mapping establishes a direct correspondence between the measurement directions and outcomes (labeled by $\vec{n}$) and a family of operators in the circuit representation parameterized by $\vec{n}$ (Fig. 2.10).

The second step is to establish a connection between $\tilde{w}(\vec{n})$ and $w(e_{\vec{n}})$ such that identical spacetime configurations in both setups have the same probability. The probability of a specific trajectory in the tensor network can be expressed as $P_{\{\vec{n}_i\}} = \prod_i w\left(e_{\vec{n}_i}\right) \left|\left| \prod_i \frac{K_{\vec{n}_i}}{\sqrt{2}} |\psi_0\rangle \right|\right|^2$, with $|\psi_0\rangle$ being the initial state / lower boundary of the toric code. By recursively defining $|\psi_t\rangle = \frac{K_{\vec{n}_t}|\psi_{t-1}\rangle}{||K_{\vec{n}_t}|\psi_{t-1}\rangle||}$ the probability can be written as

$$P_{\{\vec{n}_i\}} = \prod_t \frac{w\left(e_{\vec{n}_i}\right)}{2} ||K_{\vec{n}_t} |\psi_{t-1}\rangle\,||^2,$$

which establishes the correspondence with the trajectory probability in the circuit representation.

The $w(e_{\vec{n}})$ ensemble of measurements on the toric code is thus the same as the $\tilde{w}(\vec{n}) = \frac{w(e_{\vec{n}})}{2}$ ensemble of non-unitary circuits on Majorana fermions, where the Gaussian operations $K_{\vec{n}}$ are sampled with protocol probability $\tilde{w}(\vec{n})$ (with the restriction $\tilde{w}\left(\vec{n}\right) = \tilde{w}\left(-\vec{n}\right)$). This establishes the mapping between the measurement protocol on the toric code and a specific class of non-unitary Gaussian circuits on the virtual space. However, it is important to note that the integration domains for $\tilde{w}(\vec{n})$ and $w(e_{\vec{n}})$ are different. While $w(e_{\vec{n}})$ is defined over a hemisphere due to its axis-based definition, $\tilde{w}(\vec{n})$ is defined over the entire sphere.

2.7.4 Phase diagram of boundary MIE after general on-site measurement

Following the circuit mapping that connects different measurement protocols on the toric code with Gaussian fermionic circuits, we can readily use results of the entanglement phase diagram of non-unitary Gaussian circuits to infer the phase diagram of the MIE in the boundary state after measuring the bulk of the toric-code state on a cylinder, along any general directions. The corresponding circuit problem has been extensively studied both numerically and analytically recently; see, for example, [154, 101, 102]. Corresponding to any Gaussian circuit ensemble that satisfies the condition $\tilde{w}(\vec{n}) = \tilde{w}(-\vec{n})$, we can find the bulk measurement protocol on the toric code that realizes that Gaussian circuit, via Eq. 2.17.

These works have found the entanglement phase diagram to consist of regions of area law and critical logarithmic scaling separated by phase transitions, and we infer that the boundary MIE phase diagram also behaves similarly. This demonstrates that the MIE phase diagram we obtained by mapping the result of the specific measurement protocol (along X, Y, or Z directions) to the CPLC model is qualitatively robust to modifications of the measurement protocol to general on-site measurements. However, the specific phase boundaries and the nature of the phase transition between the area-law and critical phases vary between different ensembles, as discussed in [102].

2.8 Discussion

We explored the use of the measurement-based quantum computing (MBQC) setup for generating and manipulating quantum phases of matter. Specifically, we focused on the toric code, a topologically ordered state, and mapped the effects of random Pauli measurements to a classical loop model, allowing for an analytical understanding of measurement-induced entanglement. Additionally, we mapped general on-site measurements to Gaussian tensor networks and hybrid circuits.

We found that the entanglement pattern imprinted on the unmeasured qubits following measurement of the bulk of the toric-code ground state undergoes a phase transition that reflects the transition in the corresponding classical loop model. When a boundary chain of qubits is left unmeasured, the boundary state can have either area-law or logarithmic scaling of entanglement entropy, depending on the relative X, Y, and Z Pauli measurement frequencies. Additionally, we found that these states can also be distinguished by a spin-glass order parameter. This allowed us to devise an adaptive protocol conditioned on the

measurement results, which can steer the boundary state into a ferromagnetically ordered state, and this can be efficiently probed in experiments.

Because of the relative simplicity of bulk single-site measurements and the fact that toric-code states have already been realized in quantum hardware [204, 197], the setup described in this chapter is experimentally relevant. Our MBQC-based setup provides a way to simulate $d+1$-spacetime-dimensional hybrid circuits by one layer of local measurements on a $d+1$-dimensional entangled resource state. In quantum devices with limited coherence times, this setup may provide a promising practical route towards simulating such hybrid circuits.

This chapter illustrates how parton constructions can be leveraged in MBQC schemes, and it is worth exploring generalizations especially in higher-dimensional resource states. For example, fracton orders admit Majorana parton descriptions [98, 83], and one can consider the effect of measurements on such states. Another noteworthy example is the Levin-Wen 3D plaquette model [136], which serves as a natural generalization of the projection of a free-parton state prepared on the edges. One can analyze the effect of measurements by a very similar mapping to a loop model in three dimensions. Furthermore, considering higher-dimensional unmeasured manifolds might lead to more complex entanglement structures.

It is also interesting to investigate the relation between measurement-induced entanglement (MIE) in random bases and other aspects of wavefunction complexity. For example, MIE after measurements in a fixed basis can diagnose the sign structure in that particular basis [145], and randomizing the measurement bases may partially probe the robustness of the sign structure to local unitary transformations ("intrinsic sign structure"). There may also be connections between the universality of the resource state in MBQC and the entanglement pattern induced by measurement. For example, a similar bulk-measurement protocol on cluster states, which are universal MBQC resources, leads to a phase transition between area-law and volume-law states [147]. This can be contrasted with the toric-code case (which is not a resource for universal MBQC) in this chapter, where any subregion of the boundary state has at most logarithmic scaling of entanglement entropy.

Chapter 3

Spacetime Markov length: a diagnostic for fault tolerance via mixed-state phases

This chapter establishes a correspondence between the fault tolerance of local stabilizer codes experiencing both measurement and physical errors and the mixed-state phases of decohered resource states in one higher dimension. Drawing from recent developments in mixed-state phases of matter, this correspondence motivates a diagnostic of fault tolerance, which we refer to as the *spacetime Markov length*. This is a length scale determined by the decay of the (classical) conditional mutual information of repeated syndrome measurement outcomes in spacetime. The diagnostic is independent of the decoder, and its divergence signals the intrinsic breakdown of fault tolerance. As a byproduct, we find that decoherence may be useful for exposing transitions from higher-form symmetry-protected topological phases driven by both incoherent and coherent perturbations.

3.1 Introduction

The threshold theorem [3, 112, 114, 210] is a cornerstone of quantum fault tolerance, ensuring that quantum computations can be reliably performed if error rates are below a certain threshold. It is well known that the error threshold, in certain cases, can be related to an order-to-disorder phase transition by mapping the decoding problem to a classical statistical-mechanics model [57, 49, 88]. However, given the inherent competition between

decoherence and error correction, the threshold can be viewed instead as a transition at the level of quantum states—between distinct mixed-state phases of matter. There has been significant recent progress in understanding mixed-state phases of matter [222, 58, 85, 53, 68, 44, 14, 196, 193, 130, 245, 55, 151, 241, 152, 178, 149, 129, 244, 45, 132, 46, 131, 128, 212, 67, 230, 227, 229, 236, 81, 150, 242, 240, 148], so such a connection would allow for renewed perspectives on quantum error correction enabled by our developing understanding of quantum phases of matter.

A fundamental component of fault tolerance is quantum error correction, which often requires repeating error syndrome measurements to compensate for the fact that the measurements themselves may be faulty. In the idealistic case of perfect measurements, there has been remarkable progress in relating noisy error-correcting codes to mixed-state entanglement properties and mixed-state phases of matter [68, 14, 44, 193, 196]. The paradigmatic example of this is the toric code subjected to local Pauli noise, in which the breakdown of the quantum memory can be understood as a mixed-state phase transition induced by the noise. Note that, in this context, a mixed-state phase is taken to be an equivalence class of states, where two states are considered to be in the same mixed-state phase if there exist local Lindbladian evolutions that transform one into the other [53]. , Despite this progress, such a relation for the fully fault-tolerant setting, which includes repeated faulty measurements, has yet to be established. This necessitates taking a spacetime perspective on fault tolerance, in which physical and measurement errors appear at locations throughout spacetime [76]. This perspective is manifest in measurement-based quantum computation (MBQC), where the time dimension of a quantum error-correcting process is replaced with an extra spatial dimension [182, 21, 35]. More specifically, the spacetime quantum circuit on a system with D spatial dimentions is realized by performing measurements on a resource state in $D+1$ spatial dimensions.

In this work, we leverage ideas from MBQC to map stabilizer codes with both physical and measurement errors to the mixed-state phases of decohered resource states. One key insight gained from this mapping is that methods used to probe the mixed-state phase transition can be employed as diagnostics for fault tolerance. In particular, the conditional mutual information (CMI), defined as $I(A:C|B) = S(AB)+S(BC)-S(B)-S(ABC)$ for the annular tripartition in Fig. 3.1, has proven useful for detecting mixed-state phase transitions [193]. In particular, when the CMI decays exponentially $I(A:C|B) \sim e^{-\mathrm{dist}(A,C)/\xi}$, it defines a length scale ξ, dubbed the Markov length. The Markov length detects mixed-state phase transitions driven by local Lindbladians such as local dephasing; it is finite within each mixed-state phase but diverges at transitions [193]. The mapping motivates an analogous quantity for the stabilizer circuit, defined by the CMI of the (classical) syndrome distribution (Fig. 3.1). We show that through the mapping afforded by MBQC, the

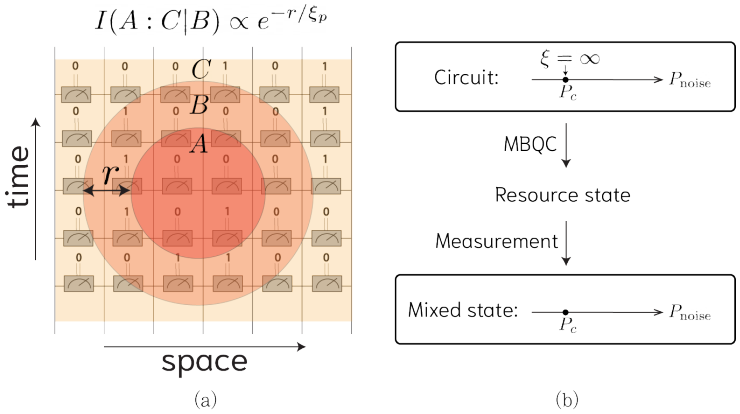


Figure 3.1: (a) A schematic of computing the CMI of a circuit with repeated syndrome measurements. The ensemble of measurement outcomes in spacetime is tri-partitioned, and its CMI is given by $I(A:C|B) = H(AB) + H(BC) - H(B) - H(ABC)$, where $H(X)$ denotes the Shannon entropy of the bits within region X. It decays with a length scale called the spacetime Markov length ξ_p. (b) The mapping of the error threshold to a mixed-state phase transition. The circuit can be mapped to a resource state for MBQC. Measurements of the resource state decohere the state to create a mixed state. Noise in the circuit maps to noise acting on the mixed state, which can drive a mixed-state phase transition occurring at the same point as the circuit threshold. The Markov length diverges at the mixed-state phase transition, signaling a similar divergence of the spacetime Markov length at the fault-tolerance threshold.

divergence of the associated length scale, which we call spacetime Markov length, signals the breakdown of fault tolerance. This correlation measure does not depend on any choice of decoder and probes the intrinsic fault-tolerance threshold.

The organization of the text is as follows. We first provide a definition of mixed-state phases in terms of local Lindbladian evolutions and in particular, emphasize that CMI is an important measure for detecting transitions [193]. We then review the foliation of syndrome-extraction circuits into resource states, demonstrating that faults in the original circuit correspond to applying noise to the resource state. Implementing the circuit corresponds to measuring all the bulk qubits of the resource state, thus leading to a decohered resource state due to both noise and measurements.

We then link the fault tolerance of the circuit to the ability to reverse the effects of noise on the decohered resource state. This allows us to identify the fault tolerance of the circuit with the condition that the decohered resource state is in the same mixed state phase as

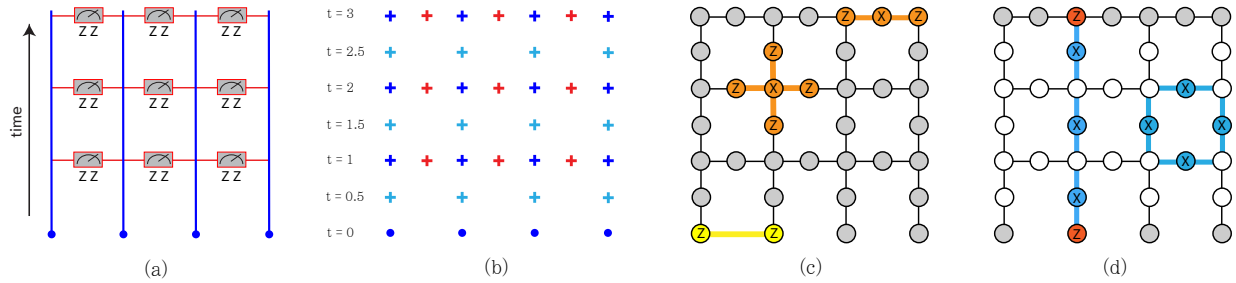


Figure 3.2: Mapping the syndrome-extraction circuit of the 1D repetition code to a 2D mixed state. (a) The syndrome-extraction circuit is composed of Z_iZ_{i+1} measurements. (b) The resource state for MBQC is defined on copies of code qubits (blue) and syndrome qubits (red). The layers are labeled by integers and half-integers. There are no syndrome qubits in the $t = 0$ layer, as the input code state is assumed to be perfect. (c) The qubits are entangled into a cluster state by applying CZ gates to neighboring qubits, with stabilizers indicated. (d) Measuring the bulk in the X basis produces a mixed state coupled to the original code on the boundaries: the resulting ensemble is stabilized by local plaquette operators and composites of the logical operators of the bulk and boundary codes.

the noiseless, albeit measurement-decohered, resource state. Subsequently, we map the CMI of the decohered mixed state back to the circuit model, in which it becomes the CMI of the classical data extracted from the spacetime history of syndrome measurements. This sets a length scale—the spacetime Markov length—which diverges at the error threshold. Finally, we discuss the experimental relevance of this diagnostic as well as applications of decoherence in probing higher-form symmetry-protected topological (SPT) phase transitions.

3.2 Mixed state phases and Markov length

We begin by reviewing a definition of mixed-state phases and the role played by the CMI. For pure ground states of local Hamiltonians, two states are in the same phase of matter if they can be connected by finite-time evolution with a local (time-dependent) Hamiltonian. For mixed states, it is natural to generalize from a local Hamiltonian to a local Lindbladian evolution. However, since such dynamics is generally irreversible, one requires the existence of both a forward and reverse evolution with local Lindbladians. That is, there needs to exist local Lindbladians $\mathcal{L}_1$ and $\mathcal{L}_2$ such that $\mathcal{T}e^{\int_0^1 \mathcal{L}_{1,2}(t)dt}[\rho_{1,2}] = \rho_{2,1}$ to conclude that ρ_1

and ρ_2 are in the same phase [53].

Given two mixed states in the same phase, finding the two-way evolutions connecting them can be challenging, and may require real-space renormalization or quantum error correction [196]. However, in many contexts, one is already given a one-way evolution $\rho_p \equiv \mathcal{T} e^{\int_0^{t_p} \mathcal{L}_1(t)dt}[\rho_1]$ (e.g. a noisy ρ_1) and the question of whether ρ_1, ρ_p are in the same phase reduces to whether the noise can be undone (is there a local Lindbladian evolution from $\rho_p \to \rho_1$). To avoid having to find a reverse connection on a case by case basis, ideally one would like a criteria for the evolving state ρ_p that guarantees the existence of a reverse evolution. In other words, it is desirable to find a mixed-state analogue of a gap for ground states of local Hamiltonians; as long as the gap remains nonzero upon perturbing the Hamiltonian, the corresponding family of ground states are guaranteed to be in the same phase [31, 155].

In Ref. [193], it was found that the Markov length ξ, i.e., the length scale at which CMI decays—plays the role of the (inverse) gap for mixed states. Given a tripartition ABC, CMI is $I(A:C|B) = I(A:BC) - I(A:B)$, a difference of two mutual informations $I(A:B) = S(A)+S(B)-S(AB)$. This representation shows that CMI quantifies how much correlation between A and its complement is not captured by its correlations with B. CMI is intimately tied to the recoverability of a state from a subsystem: when CMI is small, there exists an approximate recovery channel $\mathcal{R}_{B\to AB}$ such that $\rho_{ABC} \approx \mathcal{R}_{B\to AB}[\rho_{BC}]$. Hence, if CMI decays rapidly, then a local channel acting on A can be reversed by a channel acting on a relatively small buffer region B. Concretely, Ref. [193] used the approximate Petz theorem to argue that any local Lindbladian evolution in which the Markov length remains finite can be reversed by another local Lindbladian, where the locality of the reversal is determined by the Markov length along the evolution path. Thus, if a state evolving under a local Lindbladian has a finite Markov length, then all states along the path remain within the same mixed-state phase.

As an example of a mixed state phase transition, consider the two-dimensional toric code subject to bit-flip noise. The toric code Hamiltonian is $H_{\text{t.c.}} = -\sum_{\square} A_{\square} - \sum_{+} B_{+}$, where $A_{\square} = \prod_{i\in\square} X_i$ and $B_{+} \equiv \prod_{i\in+} Z_i$. With periodic boundary conditions, its ground state subspace is four-dimensional, encoding two logical qubits with logical operators $\tilde{X}_{v,h}, \tilde{Z}_{v,h}$ consisting of products of Pauli X and Z operators, respectively. Given perfect measurements of the syndromes $A_{\square}, B_{+}$ and uniform bit-flip noise with probability p, for $p > p_c \approx 0.11$, the logical $\tilde{Z}$ information is no longer decodable from the measurement outcomes. However, logical $\tilde{X}$ commute with the noise channel and are thus decodable for all p.

This decodability transition manifests as a mixed state transition in the following sense.

Given a ground state of the toric code $|\text{t.c.}\rangle$, the bit-flip noise converts it into a mixed state $\rho_p \equiv \mathcal{E}_p^x[|\text{t.c.}\rangle\langle\text{t.c.}|]$. Here, $\mathcal{E}_p^x[\cdot] \equiv (1-p)(\cdot) + pX(\cdot)X$ is a bit-flip channel on every site that can be represented by a continuous Lindbladian evolution $\mathcal{L}[\rho] = \sum_i \frac{1}{2}(X_i\rho X_i - \rho)$ for time $t_p = -\ln(1-2p)$. The Markov length of the mixed state ρ_p is finite everywhere except $p_c \approx 0.11$, at which it diverges [193]. Hence, the mixed state phase transition as defined by two-way local Lindbladian connectivity to the initial state ρ_0 coincides with the decodability transition.

For $p > p_c$, the fact that Markov length is finite implies that ρ_p are in the same phase as $\rho_{cl} \equiv \rho_{p=0.5} = \big(\prod \frac{1+A_\square}{2}\big)\rho_{\tilde{X}}$, which is an ensemble of loop configurations in the X basis and $\rho_{\tilde{X}}$ depends on the initial logical state. The classical ensemble with definite logical $\tilde{X}$ operators (e.g. $\tilde{X}_{v,h} = +1$) is preparable from a product state by a local Lindbladian [193] and thus in the trivial phase under the definition in [53].

However, the classical ensemble ρ_{cl} is non-trivial in several senses, which will be important for the purposes of our work. Most importantly, it is a topological classical memory with logical information given by $\tilde{X}_{v,h} = \pm 1$, and the logical states $(\prod \frac{1+A_\square}{2} \prod_{l=v,h} \frac{1\pm\tilde{X}_l}{2})$ are locally indistinguishable. This topological degeneracy is preserved under any local Lindbladian evolution in which Markov length remains finite [194]; the approximate Petz recovery for such an evolution only depends on local reduced density matrices which are the same for all topologically degenerate states. This reverse evolution thus serves as a decoder that recovers the initial state (and any logical information encoded therein) [193]. Hence, any finite-time Lindbladian evolution connecting non-trivial to trivial topological degeneracy must produce a diverging Markov length at some point [194].

This motivates the following refined definition for mixed state phases:

Definition 3.1 ([194])**.** $\rho_{1,2}$ are in the same mixed state phase if $\rho_2 = \mathcal{T}e^{\int_0^1 \mathcal{L}(t)dt}[\rho_1]$ for a local Lindbladian $\mathcal{L}$ *and* $\mathcal{T}e^{\int_0^{t'} \mathcal{L}(t)dt}[\rho_1]$ has finite Markov length for all $t' \leq 1$.

Under this refined definition, the classical state ρ_{cl} is non-trivial. Indeed, upon applying a dephasing channel $\mathcal{E}_p^z[\cdot] \equiv (1-p)(\cdot) + pZ(\cdot)Z$, the state $\rho_{cl,p} \equiv \mathcal{E}_p^z[\rho_{cl}]$ undergoes a phase transition to a trivial phase at $p_c \approx 0.11$, at which the Markov length diverges.

3.3 Mapping circuits to mixed states

Here, we describe how a circuit for syndrome extraction can be mapped to a mixed state, allowing us to use the Markov length to analyze the circuit. The first step is to map the

circuit to a resource state for MBQC. This can be accomplished using the notion of foliation, introduced in Refs. [21, 35]. For simplicity, we restrict our attention to syndrome-extraction circuits for stabilizer codes and subsystem codes in D dimensions with geometrically local checks. In this case, foliation produces a resource state in $D+1$ dimensions, which is short-range entangled in the bulk. Performing measurements on the resource state then reproduces the effects of the circuit and produces a mixed state according to the distribution of measurement outcomes. Note that, for simplicity, we consider Calderbank-Steane-Shor (CSS) codes throughout the text. We fully expect that the discussion can be extended to non-CSS codes using the generalizations in Ref. [35].

Circuit to resource state – Following Ref. [35], the resource state for MBQC is defined in a Hilbert space composed of two layers of qubits for every implementation of the circuit. This is depicted in Fig. 3.2 for the 1D repetition code. The two layers correspond to the Z-checks and X-checks of the syndrome-extraction circuit, respectively. We refer to the layers by integers and half-integers, so that the m and $m+\frac{1}{2}$ layers define the m^{th} application of the circuit, for $m \in \mathbb{Z}$. We take the initial implementation of the circuit to be the $t=0$ and $t=\frac{1}{2}$ layers, while the final implementation is $t=m_f$, for some $m_f \in \mathbb{Z}^+$. Note that we do not include a layer for $m_f+\frac{1}{2}$.[1]

Each layer is comprised of two types of qubits. First, each layer hosts a copy of the code qubits, so that the Hilbert space of the input code is reproduced in each layer. Second, there is a syndrome qubit for every check, with the placement of the syndrome qubit depending on the type of check. For the Z-type checks, we place the syndrome qubit in the integer layers, while for the X-type checks, the qubit is placed in the half-integer layers. Note that, for classical codes (with only Z-type checks), such as the repetition code, there are no syndrome qubits in the half-integer layers.

To build the resource state, we next need to specify an initial state for the code qubits in the $t=0$ layer. The choice of initial state is dependent on the context. If, for example, the circuit is being used for state preparation, then the initial state may be a product state. Here, we take the initial state to be a code state, and interpret the circuit as implementing a quantum memory. We assume that the initial code state has been prepared fault tolerantly, prior to applying the syndrome-extraction circuit.

The final step is to prepare all of the qubits (except for the code qubits in the $t=0$ layer) in the $|+\rangle$ state and couple them together with CZ gates. Specifically, we apply a CZ gate between each code qubit in the layer t and its copies in the neighboring $t-\frac{1}{2}$ and

[1] This is primarily so that, if the resource state is extended to larger t, we can take m_f to be the initial integer layer.

$t + \frac{1}{2}$ layers.[2] We also apply a CZ gate between each syndrome qubit and the code qubits of the corresponding check. Thus, in the bulk, the state is a graph state. This completes the construction of the resource state. The resource state for the 1D repetition code is shown in Fig. 3.2.

Resource state to mixed state – Now that we have defined the resource state for MBQC, we are able to finish the mapping of the syndrome-extraction circuit to a mixed state. In practice, the syndrome-extraction circuit is implemented within the MBQC formalism by measuring all of the qubits in the X basis, except for the code qubits in the final layer, that is, $t = m_f$. Given that the measurement outcomes are probabilistic, this produces a mixed state that is a classical distribution of the X basis states (except on the code qubits at $t = m_f$).

The effect of measuring the single-site X operators is equivalent to adding full-strength bit-flip noise to the resource state.[3] Therefore, letting ρ_{RS} denote the resource state, the mixed state obtained after the measurements is:

$$\rho_0 \equiv \mathcal{E}^x_{p=\frac{1}{2}}(\rho_{RS}). \tag{3.1}$$

Here, the channel $\mathcal{E}^x_{p=\frac{1}{2}}$ is the full-strength bit-flip channel, which, for each site, applies a Pauli X operator with probability $p = \frac{1}{2}$. Implicitly, $\mathcal{E}^x_{p=\frac{1}{2}}$ acts on every qubit except for the code qubits in the final layer $t = m_f$.

To gain intuition for the mixed state ρ_0, we consider the bulk state $\rho_{\mathrm{cl}} = \mathrm{Tr}_{m_f}(\rho_0)$, where we have traced out the top layer $t = m_f$. The state ρ_{cl} is a purely classical state. This is because, due to the full-strength bit-flip noise, the state ρ_{cl} commutes with every single-site X operator. Therefore, ρ_{cl} is diagonal in the X basis, i.e., it is a mixture of X basis states. For example, starting with the 1D repetition code, the state ρ_{cl} is a mixture of loop configurations, where the loops are formed by $|-\rangle$ states on the dual lattice. We argue below that, in general, the threshold of the underlying quantum code coincides with a phase transition in the bulk of the mixed state ρ_0.

[2] Here, t ranges from $0 < t < m_f$.

[3] This can be seen directly for a single qubit in the state ρ, since the post-measurement state is $\frac{1+X}{2}\rho\frac{1+X}{2} + \frac{1-X}{2}\rho\frac{1-X}{2}$, which is equal to $\frac{1}{2}\rho + \frac{1}{2}X\rho X$.

3.4 Mapping noisy circuits to noisy states

We demonstrate that noise in the circuit model corresponds to the evolution of the mixed state ρ_0 under a particular local Lindbladian. This naturally motivates the question of whether the mixed state undergoes a phase transition induced by the Lindbladian evolution. In the sections that follow, we argue that the question is equivalent to identifying whether a threshold exists for the syndrome-extraction circuit and furthermore that the mixed-state phase transition coincides with the threshold.

Consider an X or Z error occurring on a code qubit i between two consecutive implementations of the circuit, denoted by t and $t+1$. As discussed previously, each time step t in the circuit corresponds to two time steps in the resource state, specifically t and $t+\frac{1}{2}$. This mapping yields the following translation of errors:

1. **Circuit X error:** This becomes a Z error in the resource state at time step $t+\frac{1}{2}$ on the corresponding code qubit i.
2. **Circuit Z error:** This becomes a Z error in the resource state at time step $t+1$ on the corresponding code qubit i.
3. **Readout error:** This becomes a Z error on the corresponding syndrome qubit of the resource state.

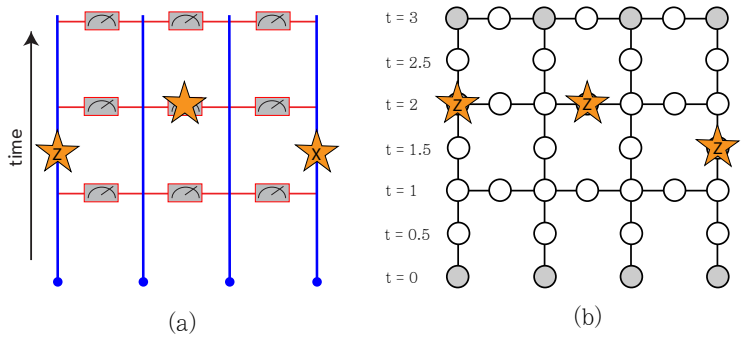


Figure 3.3: Error mapping from the syndrome-extraction circuit to the resource state for the 1D repetition code, as an example. Error locations are denoted by stars and are translated to the MBQC setting according to rules 1.-3. in the main text.

To understand this noise mapping, it is instructive to examine how qubit i is teleported from time t to $t+1$. As described in Ref. [35, 88], this is achieved by measuring this code

qubit at times t and $t + \frac{1}{2}$. After these two measurements, the state of code qubit i at time t is teleported to time step $t + 1$, though with a Pauli correction depending on the measurement results. Let us denote the two measurement results by $s_{i,t}$ and $s_{i,t+\frac{1}{2}}$, where they take the value 0 when the outcome is the $|+\rangle$ state and 1 when the outcome is the $|-\rangle$ state. These measurement results alter the teleported state at time $t+1$ by the action of the Pauli byproduct operators:

$$Z^{s_{i,t}} X^{s_{i,t+\frac{1}{2}}}. \tag{3.2}$$

This demonstrates that changing the value of the measurement result at integer time will manifest as an additional Z operator acting on the teleported state, while changing the result at half-integer time will result in an additional X operator acting on the state. The change in the measurement results can be viewed as a Z error on a code qubit, as this flips the value of s. Furthermore, a Z error on the syndrome ancilla flips the outcome of the X measurement, manifesting as a readout error. This unifies all stochastic circuit errors as correlated Z-dephasing errors in the resource state.

If these errors occur randomly in the circuit model with probability p, they will similarly manifest in the resource state with probability p. (For simplicity, we assume Pauli noise and readout error occur with the same probability, but the mapping readily generalizes.) Therefore, the resource state undergoes a dephasing channel $\mathcal{E}_p^z(\cdot) = (1-p)(\cdot) + pZ(\cdot)Z$, on the syndrome qubits and code qubits depending via the above map on the type of Pauli error. The resulting mixed state after both noise and measurement is

$$\rho_p \equiv \mathcal{E}^x_{p=\frac{1}{2}}(\mathcal{E}_p^z(\rho_{RS})) = \mathcal{E}_p^z(\rho_0), \tag{3.3}$$

as $\mathcal{E}_p^z, \mathcal{E}^x_{p=1/2}$ commute.

3.5 Fault tolerance and mixed-state phase transition

We now establish a connection between the fault tolerance of the syndrome-extraction circuit and the mixed-state phases of the measured resource state. We begin by showing that the fault tolerance of the circuit is equivalent to the fault tolerance of the corresponding MBQC protocol.

We consider a circuit or an MBQC protocol fault tolerant if they admit an error threshold for successful recovery (with a realistic error model), and they are implemented below

threshold for a sufficiently large system size. Given a noisy circuit with error rate p, we let $\tilde{\rho}_p = \sum_{\vec{S}} p_{\vec{S}} |\vec{S}\rangle\langle\vec{S}| \otimes \rho_{t,\vec{S}}$ denote the mixture of all possible syndromes $\vec{S}$, weighted by their probabilities $p_{\vec{S}}$, with the corresponding final states $\rho_{t,\vec{S}}$ of the circuit. Fault tolerance then requires the existence of a recovery map $\tilde{R}$ such that $\tilde{R}(\tilde{\rho}_p) \approx \tilde{\rho}_0$ for $p < p_c$, where p_c denotes the noise threshold and $\tilde{R}$ is independent of the encoded information.[4] Similarly, fault tolerance in the MBQC setting requires the existence of a recovery map R such that $R(\rho_p) \approx \rho_0$ for $p < p_c$.

Fault tolerance in the circuit and MBQC setting are equivalent because the existence of a recovery channel R yields a recovery channel $\tilde{R}$ for the corresponding circuit, and vice-versa. To see this, we first note that the state $\tilde{\rho}_p$ corresponds to the state ρ_p when the code qubits are post-selected in the $|+\rangle$ state within the bulk. This post-selection can be implemented through a feedback channel F, such that $F(\rho) = \tilde{\rho}$. The feedback channel measures all bulk code qubits, obtaining outcomes $\vec{C}$, and applies unitary corrections $U_{\vec{C}}$ to all syndrome and code qubits, conditioned on the outcomes. The feedback channel can be inverted ($\rho = F^{-1}(\tilde{\rho})$), where F^{-1} applies $U_{\vec{C}}^{-1}$ with probability 1/2. Therefore, the circuit and MBQC recovery maps are related by:

$$\tilde{R} = F \circ R \circ F^{-1}. \tag{3.4}$$

For further details on the feedback channel F, we refer to Appendix B.1.

This allows us to relate the threshold of the circuit to the mixed-state phase transition of $\rho_p = \mathcal{E}_p^z(\rho_0)$. These states are always one-way connected from ρ_0 by the noise channel $\mathcal{E}_p^z$, and the question of whether they are in the same phase thus hinges on the existence of a recovery R, which guarantees a circuit recovery $\tilde{R}$ as we showed.

We first prove by contradiction that the Markov length diverges at the error threshold for the circuit. Assume instead that the Markov length remains finite at p_c and within an ϵ-neighborhood around p_c. Given the noise channel evolving the density matrix ρ_p from $p = p_c - \epsilon$ to $p = p_c + \epsilon$, a finite Markov length in this region implies that a reverse Lindbladian can be constructed, evolving the density matrix from $p_c + \epsilon$ back to $p_c - \epsilon$. This reverse Lindbladian evolution, when combined with the recovery map R below the threshold, would enable recovery[5] above p_c, contradicting the assumption that p_c represents

[4] This recovery condition assumes that the fault-tolerant circuit correct all errors. In practice, there can be residual uncorrected errors that can be handled in future iterations of the circuit. Hence we disregard these residual errors by assuming that no errors occur in the final round of the circuit. Equivalently, for MBQC, this corresponds to the noise channel $\mathcal{E}_p^z$ acting on all qubits except those at the top boundary.

[5] Note that the combination of the reverse Lindbladian and a sub-threshold recovery is a valid recovery because the approximate Petz recovery of the reverse Lindbladian depends solely on local reduced density matrices, which are independent of the encoded logical information of the underlying topological code.

the true threshold.

While this argument is sufficient to show that the Markov length diverges at p_c, it does not exclude the possibility of divergence at multiple points. However, we expect that this generically does not happen, based on mapping the problem of recovery to statistical-mechanics models [57]. In this mapping, the (dis)ordered phase of the stat mech model corresponds to the (un)recoverable regime. The CMI corresponds to the free energy cost of a point defect in the stat mech model [193], and thus, the Markov length can be identified with the correlation length—assuming that there is only one diverging length scale in the vicinity of the critical point of the stat mech model. As a consequence, the divergence of the Markov length coincides with the recoverability transition.

For simplicity, we have considered incoherent noise on the syndrome extraction circuit. In the case of non-stochastic errors, such as coherent errors, the mapping of the errors to the resource state is more subtle. In Appendix B.2, we discuss the subtleties and show that for the repetition code in particular, coherent noise on physical qubits in the circuit model is mapped to incoherent noise on the resource state – allowing for a connection to a mixed-state phase transition. We also note in Appendix B.3 that coherent noise on the MBQC resource state is equivalent to Z-dephasing errors on the X-dephased resource state. This implies that the threshold of MBQC to coherent noise corresponds to a mixed-state phase transition of the X-dephased resource state under incoherent noise.

3.6 Classical memory in the bulk

Here, we provide another perspective on the breakdown of the fault tolerance of the circuit. In particular, we show that it coincides with a transition in the recoverability of a classical memory in the bulk of the decohered resource state. To specify the classical memory, we first identify two important classes of stabilizers of the resource state, which we refer to as *detector cells* and $L'BL$. We note that, later in the text, the detector cells are important for establishing that the spacetime Markov length of the circuit measurement outcomes gives a diagnostic of the fault tolerance.

Detector cells – The first type of stabilizers, referred to as detector cells following Refs. [107, 169], capture the redundancy inherent to the circuit. Without errors, consecutive measurements of a stabilizer yield consistent syndromes. For instance, in the repetition code, the measurements of the stabilizer Z_iZ_{i+1} at times t and $t+1$ should agree in the absence of errors. In MBQC, this process is equivalent to first measuring the X operator

on the syndrome qubit associated with the stabilizer. Then, measuring X_i and X_{i+1} on the code qubits in the intermediate layer $(t+\frac{1}{2})$ teleports the stabilizer from t to $t+1$, possibly with a sign change X_iX_{i+1}. Finally, measuring X on the syndrome qubit at $t+1$ yields a second syndrome measurement, influenced solely by the teleportation-induced sign change. Thus, assuming no errors, the product of these X measurements should yield $+1$. This detector cell is depicted in Fig. 3.2(d) as a plaquette of X operators.

In general, for a stabilizer S measured in the circuit, the corresponding detector cell is:

$$D_S = \left(\prod_{i\in \mathrm{Ch}(S)} X_i^{(t+1)}\right)\left(\prod_{i\in \mathrm{Sup}(S)} X_i^{(t+\frac{1}{2})}\right)\left(\prod_{i\in \mathrm{Ch}(S)} X_i^{(t)}\right). \tag{3.5}$$

Here, $\mathrm{Ch}(S)$ is a choice of syndrome qubits such that the product of the corresponding checks is S,[6] and $\mathrm{Sup}(S)$ is the set of code qubits in the support of S. Measuring X-operators on $\mathrm{Sup}(S)$ in the intermediate layer $t+\frac{1}{2}$ teleports the stabilizer S from layer t to $t+1$, with the X-measurements at $t+1$ providing a second syndrome consistent with that at t.[7]

$L'BL$ stabilizers – The second type of stabilizer arises from the temporal consistency of logical operators. Without noise, the logical information encoded in the code at different time instances should agree. As an example, consider two instances of a repetition code logical operator, $Z_i^{(m_f)}$ and $Z_i^{(0)}$, acting on the code qubits in the final and initial layers of the circuit. The temporal consistency of this logical is reflected in the resource state through a stabilizer of the following form:

$$L^Z_{m_f} B_{L^Z} L^Z_0 = Z_i^{(m_f)} \left(\prod_{k=0}^{m_f-1} X_i^{(k+\frac{1}{2})}\right) Z_i^{(0)},$$

Here, the logical operators $L^Z_{m_f}$ and L^Z_0 are linked via the operator B_{L^Z}, which acts within the bulk of the state. This stabilizer is illustrated in Fig. 3.2(d), where the supports of $L^Z_{m_f}$ and L^Z_0 are highlighted in red, and the support of B_{L^Z} is depicted in blue.

[6] Note that the choice of syndrome qubits is for subsystem codes in particular, where different choices of gauge operators may multiply to the same stabilizer.

[7] Detector cells can also correspond to local redundancy checks within a single layer. For example, in the 3D toric code, there are detector cells represented by the product of plaquette stabilizers around a cube. However, we do not explicitly discuss these types of detector cells here, for simplicity.

In general, the resource state has stabilizers of the form $L_{m_f} B_L L_0$, with B_L given by

$$B_{L^Z} = \prod_{k=0}^{m_f-1} \prod_{i \in \mathrm{Sup}(L^Z)} X_i^{(k+\frac{1}{2})}, \qquad B_{L^X} = \prod_{k=0}^{m_f-1} \prod_{i \in \mathrm{Sup}(L^X)} X_i^{(k)}, \tag{3.6}$$

depending on whether the logicals L_0, L_{m_f} are products of Pauli Z or X operators. We refer to the operators of the form $L_{m_f} B_L L_0$ as $L'BL$ stabilizers, for simplicity.

Classical memory in the bulk – We now explain how the B_L and D_S operators together define a classical memory. In this memory, the D_S operators act as the check operators, while the B_L operators represent the logical bits. Both D_S and B_L are products of X-type operators and thus commute. Moreover, the B_L operators are independent of the D_S operators because logical bits represented by B_L can be flipped by conjugating them with L_m or $L_{m-\frac{1}{2}}$,[8] while the D_S operators remain invariant (a logical error does not alter the values of detector cells).

In the repetition code example, the bulk code has plaquettes of X-type operators serving as checks and open strings of X-type operators connecting one boundary to another as logical operators, as depicted in Fig. 3.2. The code distance scales linearly with the system size, as non-contractible closed loops of Z-errors can alter the logical bit without being detected by the checks. Note that for the 1D repetition code, there is only one bit encoded, corresponding to the operator B_{L^Z}. In contrast, there is no bit corresponding to the operator B_{L^X}, as the repetition code is designed to protect against X-type noise only.

Error correction and fault tolerance – We now analyze how the error correction properties of this classical code relate to the fault tolerance of the circuit. Let us begin with the noiseless scenario. In this case, transmitting a logical state from one boundary to another in MBQC can be understood as a two-step process. First, measurements in the bulk establish entanglement between the top and bottom boundaries. Then, measuring the bottom boundary teleports the information to the top boundary, upon applying feedback depending on all measurement outcomes.

The $L'BL$ stabilizer enables this teleportation process. The initial state satisfies $L'BL = 1$, and measuring the bulk corresponds to measuring the B operators. Consequently, the

[8] Here, L_m and $L_{m-\frac{1}{2}}$ are the logical operators of the underlying code defined within the m and $m-\frac{1}{2}$ layers of the resource state, respectively.

state on the boundaries is stabilized by $L'L$, taking value B. Successfully teleporting the logical state from one boundary to the other requires feedback depending on the measured value of B.

In the presence of noise, the errors need to be removed while preserving $L'BL$; otherwise the feedback is wrong, and we have committed a logical fault. The condition for removing errors without changing a logical operator B is precisely the condition for the classical memory to persist in the presence of noise. Thus, retention of classical memory implies the ability to teleport the logical state from the initial boundary to the final boundary, which in turn implies preservation of logical information in the original circuit.

As the simplest explicit example, the threshold for the 1D repetition code with probability p for both bit-flip and readout error maps to the threshold for the 2D classical memory defined by $\rho_{cl} = \left(\prod \frac{1+A_\square}{2}\right)\rho_{\tilde{X}}$ subject to dephasing noise $\mathcal{E}_p^z$. We note that [242] offers a complementary approach to our mixed state phase perspective, using strange correlators of the noisy cluster state to diagnose its ability to transmit information from one boundary to the other.

Higher-Form Symmetries – In some cases, the detector cells generate higher-form symmetries, i.e., for which the symmetry operators are supported on closed submanifolds, as described in Ref. [176]. In particular, when the stabilizers generate a q-form symmetry, the detector cells then generate a q-form symmetry in one dimension higher. For example, the stabilizers of the toric code generate a $\mathbb{Z}_2 \times \mathbb{Z}_2$ 1-form symmetry. Correspondingly, the detector cells of the foliated state also generate a $\mathbb{Z}_2 \times \mathbb{Z}_2$ 1-form symmetry—this is the well-known $\mathbb{Z}_2 \times \mathbb{Z}_2$ 1-form symmetry of the 3D cluster state [184]. Thus, in this case, the redundancy of the stabilizer measurements produces a higher-form symmetry of the resource state.

This contrasts with cases such as the Bacon-Shor code, which is not fault tolerant. In that case, the detector cells of the resource state are rigid, in the sense that they are linear subsystem symmetries of the resource state. We postulate that, if the checks are local, then higher-form symmetries of the resource state are necessary for fault tolerance.

In the case that the stabilizers of the underlying code generate an anomalous q-form symmetry, then the stabilizers of the form $L'BL$ are an indication that the resource state belongs to a nontrivial q-form SPT phase. The $L'BL$ operators imply that the symmetry defects of the q-form symmetry (created by B) are dressed with q-form symmetry charges, given by L' and L. This is because the operator L_m ($L_{m-\frac{1}{2}}$) defined by the logical operator L^Z (L^X) fails to commute with B_{L^X} (B_{L^Z}) when they have overlapping support and L^Z and L^X are conjugate logical operators.

For example, the stabilizers of the toric code generate an anomalous $\mathbb{Z}_2 \times \mathbb{Z}_2$ 1-form symmetry. The resource state is known to belong to a nontrivial 1-form SPT phase, as seen by the $L'BL$ stabilizers. We refer to Refs. [184, 168] for further details.

3.7 Spacetime Markov length

We employ the Markov-length as a diagnostic for a mixed-state phase transition to show that the CMI of the classical data extracted from a noisy circuit serves as an indicator of the error threshold. We begin by showing that the CMI of a decohered resource state depends exclusively on the entropy of specific detector cells. We then demonstrate a parallel result for the circuit itself, where the CMI of classical syndromes depends solely on the entropy of the corresponding classical bits associated with the detector cells. Finally, we prove that under the noise mapping, the entropy of the resource state's detector cells matches the entropy of their circuit counterparts. This correspondence establishes an equality between the CMI of the decohered resource state and that of the classical data, thereby confirming spacetime Markov length as a robust indicator of the error threshold.

We begin by examining a decohered resource state derived from the correspondence with a given quantum circuit. To quantify the CMI of this state, we compute the reduced density matrix for a given subregion A in the bulk of the decohered resource state and analyze its entropy. Decohering the system in the X-measurement basis effectively removes all initial stabilizers of the state except the ones that commute with X-measurements. These stabilizers are composed of Pauli-X operators and are the detector cells. Consequently, the reduced density matrix post-measurement (with discarded outcomes) assumes the form:

$$\rho_A = \frac{1}{N} \prod_k \frac{1 + D_k}{2}, \tag{3.7}$$

where N is a normalization factor, and D_k are stabilizers corresponding to detector cells inside the region.[9]

Introducing a noise channel with strength p to this reduced density matrix modifies it to:

$$\rho_{A,p} = \frac{1}{N} \sum_{m_k} P_{m_k} \prod_k \frac{1 + m_k D_k}{2}, \tag{3.8}$$

[9] In general, non-local symmetries may exist within a given region, particularly for regions that are not simply connected; however, we still refer to these as detector cells.

where P_{m_k} represents the probability distribution over the sign of detector cells (m_k) in the resource state. Computing the von Neumann entropy of this density matrix gives:

$$S(\rho_{A,p}) = H(m) + |A| - |D| \tag{3.9}$$

This has two contributions: the first is the Shannon entropy of the probability distribution of the sign distribution, $H(m) = \sum P_{m_k} \log P_{m_k}$, and the second is a term from the normalization factor N, which depends on the number of qubits inside the region ($|A|$) and the number of detector cells within it ($|D|$).

Next, we examine the corresponding region A in the circuit. In a noise-free scenario, the extracted syndromes have deterministic components that correspond to detector cells in the resource state. For instance, consider two syndromes s_1 and s_2 that are outcomes of repeated measurements of a given stabilizer. From the consistency of the result in the noise-free scenario, we know $s_1 = s_2$ which means $d_1 = s_1 s_2$ is a deterministic variable. Given a bit string s_i from region A, the syndromes can be expressed in terms of deterministic components d_i and a basis for non-deterministic components n_i. The non-deterministic syndromes remain completely random, even under stochastic errors, as seen from the stabilizer formalism. Thus, the entropy of s_i includes two contributions: one from the deterministic components d_i and another from the fully random components n_i. Stochastic errors turn the deterministic components d_i into random variables, with probability distribution $P(d_1, d_2, \ldots, d_k)$. This yields the total entropy:

$$H(s) = H(d) + |s_i| - |d_i|, \tag{3.10}$$

where $H(d)$ represents the entropy of the d_i values, and the second two terms account for the $|n_i|$ fully random variables within the region.

The correspondence between detector cells in the resource state and deterministic components in the circuit implies $|d_i| = |D|$, and the noise mapping ensures that any error changing the sign of a detector cell also changes the sign of the corresponding deterministic component. This shows that the probability distribution of the sign of detector cells P_m matches that of the deterministic components P_d, so $H(m) = H(d)$. Because the volume law terms $|s_i|$ and $|A|$ do not contribute to CMI, this establishes the equality of CMI between the circuit and MBQC.

3.8 Discussion

We have shown that the breakdown of fault tolerance in local stabilizer quantum memories can be mapped to mixed state transitions of classical memories in one higher dimension. This mixed-state perspective offers the novel diagnostic of the spacetime Markov length, which can be computed directly from the syndrome distribution of the circuit. This diagnostic does not depend on any decoder and thus probes the intrinsic threshold transition.

As repetition and surface codes below the noise threshold have been realized recently on quantum hardware [4], it would be very interesting to extract the spacetime Markov length from experimental syndrome data as a complementary probe of the proximity to the threshold. Real-time decoders acting in parallel on space-time blocks (and then finding a globally consistent solution) have been demonstrated [4]. The spacetime Markov length scale, in particular, may be useful to estimate the minimum block sizes for such decoders.

This approach can also be applied to more general circuits that implement logical gates through code deformations, such as lattice surgery. These circuits define resource states where the logical gates are encoded into the geometry of the cluster state [35]. Further research could also investigate mixed-state phase transitions within the formalism of ZX calculus [23], which would generalize our analysis to inherently dynamical quantum-error correcting codes, such as Floquet codes [86]. Another promising direction is the exploration of spacetime circuit codes, where we anticipate that the outcome code [56] could serve a similar role as the classical memory in our analysis. In this context, error correction for the classical memory would be linked to fault tolerance, and the spacetime Markov length could once again be used to detect the threshold.

While this work focused on fault-tolerant stabilizer codes (and corresponding Clifford operations), it would be interesting to extend this approach to non-Abelian topological memories or operations beyond the stabilizer paradigm. The general case requires adaptive feedback from syndrome measurement, meaning that the mapping of a fault-tolerant circuit to a noisy resource state would be more complex. However, in this context, the spacetime Markov length may still be relevant. Fault tolerance with measurement errors relies on the redundancy of the syndromes, and thus it is plausible that the classical syndrome state must generally be non-trivial and spacetime Markov length may detect its transition to a trivial state.

Application to higher-form SPT transitions: We note that the measurements which decohere the resource state into a classical ensemble are essential for defining a mixed state phase transition. Without measurements, the resource state with periodic boundary conditions does *not*, in fact, undergo a phase transition when noise is applied. This is

because a cluster state subject to local dephasing noise is equivalent to a thermal Gibbs state of the cluster-state Hamiltonian [181, 45], which has zero CMI [36].

The measurement decoherence which exposes the mixed state phase transition may be more generally useful in exposing "hidden" transitions of SPT phases protected by higher-form symmetries under both incoherent and coherent perturbations. For example, in Ref. [220], there are instances in which the higher-form symmetry protecting an SPT is explicitly broken, leading to no bulk phase transition but which still exhibits a boundary transition. The effect of the measurement decoherence in the symmetry-charge basis is to effectively introduce a boundary after spacetime rotation [14], thus exposing the hidden transition.

Chapter 4

Critical non-equilibrium phases from noisy topological memories

This chapter demonstrates the existence of an extended non-equilibrium critical phase, characterized by sub-exponential decay of conditional mutual information (CMI), in the surface code subject to heralded random Pauli measurement channels. By mapping the resulting mixed state to the ensemble of completely packed loops on a square lattice, we relate the extended phase to the Goldstone phase of the loop model. In particular, CMI is controlled by the characteristic length scale of loops, and we use analytic results for the latter to establish polylogarithmic decay of CMI in the critical phase. We find that the critical phase retains partial logical information that can be recovered by a global decoder, but not by any quasi-local decoder. To demonstrate this, we introduce a diagnostic called punctured coherent information which provides a necessary condition for quasi-local decoding.

4.1 Introduction

Error correction involves the encoding of information into long-range correlations of many-body states [209], resulting in non-trivial phases of matter such as topological order [232, 112]. Topological codes such as toric code exhibit noise thresholds above which encoded information is no longer decodable [57]. Interpreting this decodability transition as a phase transition has led to fundamental developments in generalizing the notion of a phase to mixed-states (arising from decoherence) [14, 68, 226, 53, 178, 44, 46, 193, 162, 212, 67,

191, 240, 151, 131, 130, 133, 189, 225, 146]. Unlike equilibrium critical points in which the correlation length of local observables diverges, such mixed-state phase transitions induced by noise require fundamentally distinct probes.

In particular, conditional mutual information (CMI) provides a diagnostic of such transitions and governs the local recoverability of a state against local noise. For a tripartition $A{:}B{:}C$ of a state ρ, the CMI is

$$I(A:C|B)_\rho = S(AB) + S(BC) - S(B) - S(ABC), \tag{4.1}$$

where $S(\cdot)$ denotes the von Neumann entropy. Consider an annular tripartition such that A and C are completely separated by B of characteristic length r (Fig. 4.1). If $I(A:C|B)_\rho \sim \exp(-r/\xi_{\mathrm{M}})$, then the state is said to have finite *Markov length* ξ_{M}. Gapped ground states [87] and finite-temperature Gibbs states of local Hamiltonians [43, 119, 104] all have finite Markov length. At thermal (equilibrium) phase transitions, conventional correlation lengths diverge, but the Markov length is finite. Quantum phase transitions in ground states, on the other hand, are characterized by the divergence of both the correlation and Markov lengths, which coincide for pure states. In contrast, at mixed-state phase transitions of topological memories subject to local noise [57, 68, 14], the correlation length remains finite whereas the Markov length diverges [193], a fundamentally non-equilibrium phenomenon. In that setting, finite Markov length controls the quasi-local recoverability of encoded information: when the noise is generated by a short time evolution with a local Lindbladian, as long as the Markov length is finite, the effect of local noise can be undone by a quasi-local recovery channel [193].

In equilibrium, there are not only isolated critical points but also extended critical phases such as Luttinger liquids, metallic phases, or the low-temperature Berezinskii–Kosterlitz–Thouless (BKT) phase, which are described by a manifold of renormalization group fixed points [186]. Can there be *non-equilibrium* critical phases? A uniquely non-equilibrium feature for such a critical phase is sub-exponential decay of CMI despite short-ranged correlations in an extended parameter space. Given such a phase in the context of error correction and topological codes, what are the consequences for decodability?

In this work, we realize a non-equilibrium critical phase in an ensemble obtained from a topological code subject to heralded noise. We show that this ensemble has sub-exponential decay of CMI in a finite parameter range despite short-range correlations, and we relate this regime to the Goldstone phase of a loop model with crossings [160].

We then analyze the implications for error correction, finding that in this critical phase, some logical information is globally but not quasi-locally recoverable. For this purpose, we introduce a new diagnostic called punctured coherent information, which provides a

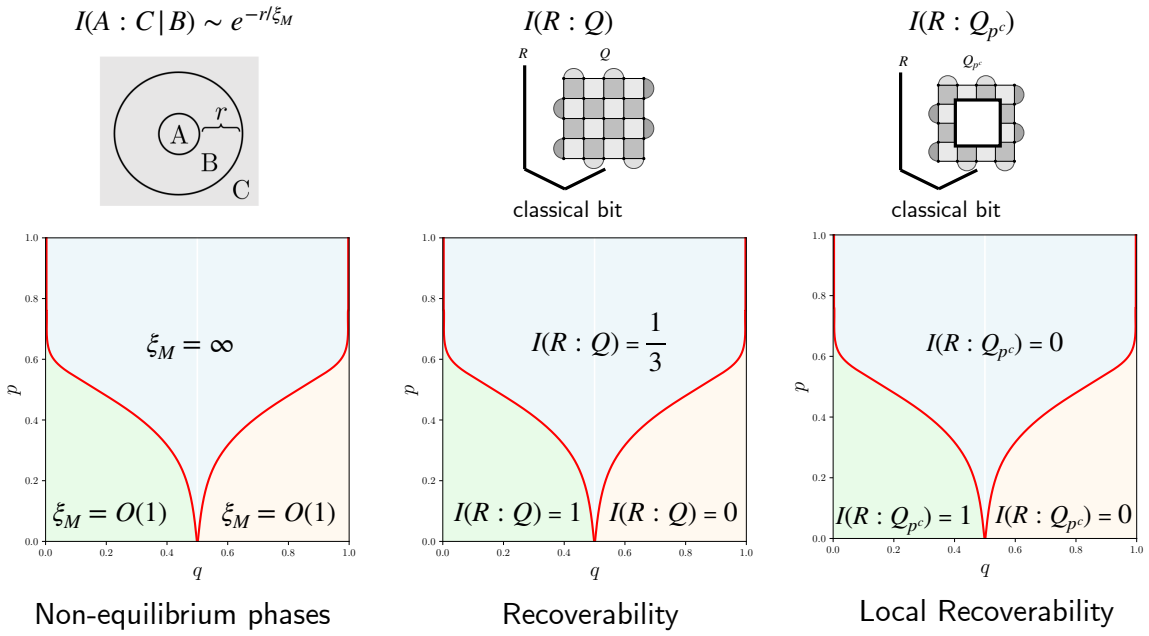


Figure 4.1: The surface code subject to heralded complete Pauli dephasing has a phase diagram described by the completely packed loop model on a square lattice, and the parameters $0 \leq p, q \leq 1$ control the relative probabilities of X, Y, and Z Paulis in the dephasing channel. There are two topologically distinct short-loop phases which exhibit finite Markov length, and an extended Goldstone phase with system-spanning loops, which exhibits diverging Markov length. We also compute the persistence of a classical logical memory using the mutual information $I(R{:}Q)$ between a classical reference bit R and the dephased code Q. In our setting, this differs from the coherent information $I_c(R\rangle Q) \coloneqq S(Q) - S(RQ)$ only by the constant $S(R) = 1$. Furthermore, we propose a punctured version of the mutual information $I(R{:}Q_{p^c})$ between the reference R and part of the code Q as a probe for local recoverability of the encoded information. Studying the interplay of phases, memory, and recoverability of encoded information in each of these phases, we find that the Goldstone phase is characterized by diverging Markov length and partial classical memory that is globally but not locally recoverable. In contrast, in the short-loop phase (shaded green), X-type logical information persists completely, and we show that it can be recovered using a quasi-local channel.

condition for quasi-local decodability, and illustrate its utility in the critical phase. We also propose an explicit quasi-local decoder in the non-critical phase retaining logical memory.

We note that recent work has shown that noisy fractional quantum Hall states [228] and decohered non-Abelian topological orders [190, 188] can exhibit an extended critical regime of algebraically decaying Renyi correlators, whereas our focus is on establishing a regime of sub-exponential CMI decay and global versus quasi-local decodability. Prior works [130, 245, 149, 243] have studied critical mixed-states arising from decohering quantum critical pure states, but these have algebraic decay of both CMI and correlations. Finally, we note that while the present work was completed, [224], which defines "information critical phases" also based on diverging Markov length, appeared.

We consider the surface code subjected to a heralded single-qubit Pauli dephasing channel, where a local three-level "flag" heralds which of the Pauli X, Y, Z bases the physical qubits were dephased in. Equivalently, the channel can be understood as a non-selective local Pauli measurement channel in the surface code. The resulting state is classical (diagonal in a product basis) but can nonetheless retain a classical topological memory.

We map this classical ensemble to a completely packed loop model with crossings (CPLC). In this representation each measurement pattern corresponds to a configuration of loops on the lattice, and each loop imposes a $\mathbb{Z}_2$ constraint on the measurement outcomes. The classical CMI of the ensemble can be expressed directly in terms of the number and structure of these loop constraints.

The main results of this chapter are:

1. *Critical mixed-state phase and diverging Markov length.*

   We map CMI of the classical ensemble to properties of the CPLC model. The CMI is given by the number of independent loop-parity constraints that connect A and C when conditioning on B. Thus, the CMI in certain tripartitions $A{:}B{:}C$ is controlled by loop spanning numbers and multi-leg (watermelon) correlators in the CPLC.

   The CPLC has a well-understood phase diagram featuring two short-loop phases and an extended Goldstone critical phase. In the Goldstone phase, the spanning number between opposite boundaries grows logarithmically with system size, while multi-leg correlators have polylogarithmic decay; in the short-loop phases these quantities decay exponentially. We show, analytically and numerically, that the CMI inherits this behavior: for a partition as shown in Fig. 4.1 left panel, the CMI decays exponentially with separation in the short loop phase (implying a finite Markov length), while it has polylogarithmic decay in the Goldstone phase (divergent Markov length).

2. *Classical memory and (non-)local recoverability.*

   Finally, we interpret the classical ensemble as a noisy classical memory. After complete X dephasing, the surface code stores a single logical classical bit. We couple this logical bit to a classical reference and track the mutual information between the reference and the noisy code. We introduce a "punctured" version of this mutual information—obtained by first tracing out a finite fraction of bulk region of the code—which probes whether decoding can be achieved quasi-locally.

   In the Goldstone phase we find that the input–output mutual information $I(R{:}Q)$ remains nonzero, showing that the logical bit is retained, but that the punctured mutual information $I(R{:}Q_{p^c})$ vanishes for any extensive bulk puncture. We establish general constraints relating local recovery maps and punctured coherent information and use them to conclude that in the Goldstone phase no quasi-local decoder can exist. Equivalently, recovery requires a decoder of intrinsically high circuit complexity and cannot be achieved by any quasi-local decoding circuit.

   On the other hand, the phase diagram also contains a logical-preserving short-loop phase in which the encoded bit is not only retained but also quasi-locally recoverable. In this regime we construct an explicit decoder: a recovery channel implementable by a circuit of polylog(L) range and depth, with failure probability vanishing polynomially as the system size $L \to \infty$. (In the other short-loop phase, the logical information is erased.)

The remainder of this chapter is organized as follows. In Sec. 4.2 we review the surface code and its parton construction, and define the heralded Pauli dephasing channel and the resulting classical ensemble. In Sec. 4.3 we relate the CMI of this ensemble to loop-counting in the CPLC and establish the existence of a critical mixed-state phase with divergent Markov length. Section 4.4 discusses the implications for classical memories, introducing punctured coherent information and constructing an explicit log-local decoder in the short loop regime. We conclude in Sec. 4.5 with a brief outlook.

4.2 Setup

4.2.1 Surface code and Majorana representation

We work with the rotated surface code on a square lattice with qubits on the vertices; see Fig. 4.2. Each plaquette p supports either an X-type stabilizer $S_p^X = \prod_{i\in p} X_i$ or a Z-type

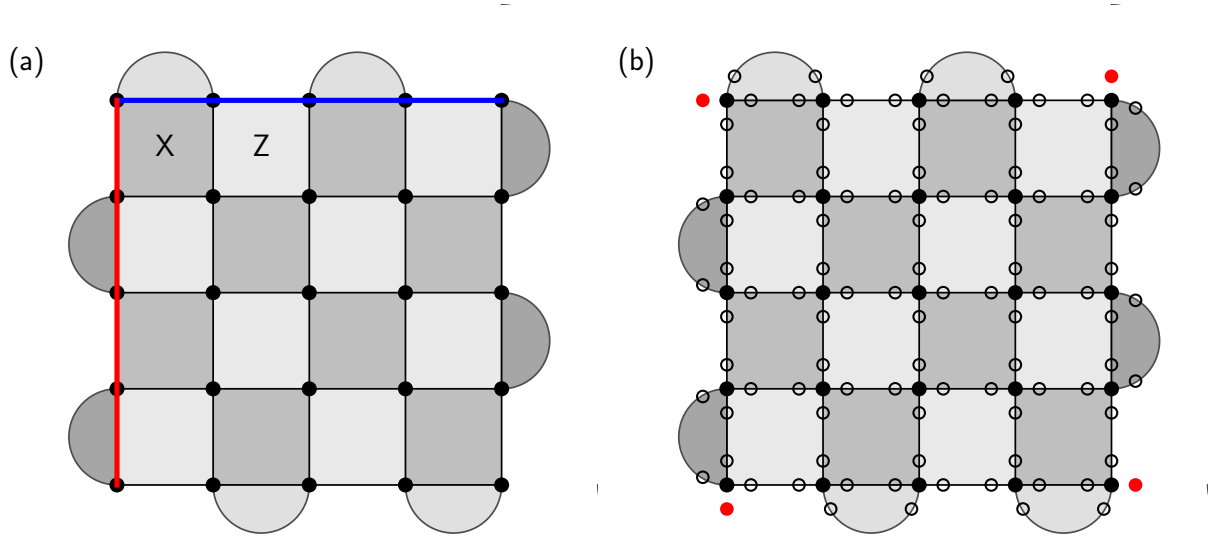


Figure 4.2: Left: A $d = 5$ rotated surface code with X-type (black) and Z-type (white) plaquette stabilizers. Logical operators $\bar{Z}$ (red) and $\bar{X}$ (blue) run along boundaries. Right: The Majorana–qubit tensor-network representation, with 4 virtual Majorana modes (unfilled circles) attached to each physical qubit leg (filled circles). Neighboring virtual Majoranas are projected into even-parity dimers along edges (see text), while the 4 unpaired corner modes (red circles) encode the logical qubit degrees of freedom.

stabilizer $S_p^Z = \prod_{i\in p} Z_i$, and the boundary is chosen so that strings of X and Z operators along the edges implement logical $\bar{X}$ and $\bar{Z}$. This is the rotated surface-code patch [22], closely related to the planar surface codes [112].

For our purposes it is convenient to represent the surface-code ground state as a projected fermionic state. This construction is based on the parton representation [233, 111, 28] of a qubit at each lattice site i in terms of four *virtual* Majorana modes $\gamma^i_{\rm up}$, $\gamma^i_{\rm down}$, $\gamma^i_{\rm left}$, and $\gamma^i_{\rm right}$.

We first define a local 5 legged tensor by specifying a pure state on the Hilbert space of *(one physical qubit)* $\otimes$ *(four virtual Majoranas)*. A convenient choice is to take the local tensor to be the projector onto the unique common $+1$ eigenspace of three commuting stabilizers: a purely fermionic parity constraint

$$S^i_1 \;=\; -\,\gamma^i_{\rm up}\gamma^i_{\rm down}\gamma^i_{\rm left}\gamma^i_{\rm right}, \tag{4.2}$$

and two mixed qubit–fermion stabilizers tying physical Pauli operators to Majorana bilinears,

$$S_2^i \;=\; X^i\,(i\,\gamma_{\rm up}^i\gamma_{\rm right}^i), \qquad S_3^i \;=\; Z^i\,(i\,\gamma_{\rm up}^i\gamma_{\rm left}^i). \tag{4.3}$$

We choose the above convention for the A sublattice, and on the B sublattice we apply a Hadamard gate H to the physical qubit, which swaps $X^i \leftrightarrow Z^i$ while leaving the Majoranas unchanged.

To construct the global tensor network, we contract the virtual Majorana legs along lattice edges by projecting neighboring modes into *dimers*. If the right leg of site i is joined to the left leg of site j, we insert the even-parity (dimer) projector

$$P_{ij}^{(\eta_{ij})} \;=\; \frac{1}{2}\Big(1+\eta_{ij}\,i\,\gamma_{\rm right}^i\gamma_{\rm left}^j\Big), \tag{4.4}$$

and similarly along all other edges. Here $\eta_{ij} = \pm 1$ is a choice of sign (equivalently, an orientation convention) for that bond. After inserting all such bond projectors and contracting all virtual Majorana legs, only the physical qubit legs remain open; the resulting state on the physical legs is the surface-code state defined by this construction.

For a fixed choice of local tensors and a fixed sign pattern $\{\eta_{ij}\}$, the contraction defines a unique physical state (up to overall normalization). However, different sign choices in the dimer projections can generate surface-code stabilizer realizations that differ by signs of plaquette operators (i.e., which plaquettes are constrained to have eigenvalue $+1$ versus -1 in the resulting stabilizer state). In particular, to obtain the standard rotated surface code in which all plaquette stabilizers are $+1$, one must choose a compatible sign/orientation convention for the dimer projectors. We follow the convention described in Ref. [28], which fixes the $\{\eta_{ij}\}$ (or equivalently an edge orientation) so that the induced plaquette signs match the usual surface-code stabilizers.

Along the outer boundary, four Majorana legs remain uncontracted. These boundary Majoranas encode the logical degrees of freedom: different ways of pairing them correspond to different logical bases, and in particular realize the logical Pauli operators $\bar{X}, \bar{Y}, \bar{Z}$ as nonlocal bilinears of boundary modes (possibly dressed by boundary dimer variables, depending on convention). The topological information of the code is thus carried by the boundary Majoranas together with the bulk pattern of dimer projections.

This parton (Majorana) construction is particularly useful for tracking the effect of single-qubit Pauli measurements. In the tensor-network picture, a Pauli measurement acts only on the physical leg of a local tensor; fixing that leg (in the appropriate Pauli eigenbasis) is equivalent to imposing a local constraint on the four virtual Majorana legs, thereby selecting one of a small set of allowed fermionic pairings ("rewirings") inside the tensor. The

virtual contractions between neighboring sites are unchanged, so a measurement pattern updates the global state simply by replacing the affected site tensors with new purely fermionic tensors implementing the corresponding local Majorana rewiring. For example, on the sublattice B (which includes the second qubit on the topmost row in Fig. 4.2), an X-measurement on site i with outcome $m = \pm 1$ (projection $\Pi_m^{(X)} = \frac{1}{2}(\mathbb{I}+mX^i)$) replaces X^i by m in the mixed stabilizer $S_2^i = X^i(i\gamma_{\text{up}}^i\gamma_{\text{right}}^i)$, yielding the Majorana constraint $i\gamma_{\text{up}}^i\gamma_{\text{right}}^i = m$; together with the parity $S_1^i = -\gamma_{\text{up}}^i\gamma_{\text{down}}^i\gamma_{\text{left}}^i\gamma_{\text{right}}^i = +1$, this implies $i\gamma_{\text{down}}^i\gamma_{\text{left}}^i = m$. Thus the physical leg is removed and the site reduces to a Gaussian tensor pairing the virtual Majoranas into dimers $(\text{up}, \text{right})$ and $(\text{down}, \text{left})$. The same logic applies for Y and Z measurements: fixing the physical leg replaces the relevant mixed qubit–Majorana stabilizer(s) by purely Majorana bilinears, selecting a definite local pairing pattern, which are listed in Fig. 4.3a. In the rest of the chapter we exploit this simplification: bulk operations acting on physical qubits (such as the heralded Pauli dephasing channel) reduce to local tensor replacements that implement these rewiring rules, while the pattern of virtual dimer projections stays fixed; contracting the resulting purely fermionic network yields either constrained loop ensembles (in the fully dephased classical description) or effective boundary states defined on the uncontracted Majoranas, providing a direct bridge between the properties of the ensemble and the loop model.

4.2.2 Heralded Pauli dephasing channel and the ensemble

We now define the heralded noise model and the resulting classical ensemble. Heralded erasure channels, where the location of erasure or loss event is known have been extensively studied [77, 113, 215]. In this work, we introduce a heralded dephasing channel, defined as follows.

On each site we introduce a three-level "flag" that indicates the local Pauli operator on which the dephasing is performed. The physical Hilbert space is $\mathcal{H}_{\text{s}} \cong (\mathbb{C}^2)^{\otimes n}$, and the flag space is $\mathcal{H}_{\text{f}} \cong \left(\text{span}\{|0\rangle, |-1\rangle, |+1\rangle\}\right)^{\otimes n}$. On each physical qubit of the surface code, we apply a local CPTP map $\mathcal{N}_{p,q}$, parametrized by two probabilities $p, q \in [0, 1]$,

$$\begin{aligned}\mathcal{N}_{p,q}(\rho) =&(1-p)(1-q)\mathcal{D}_X(\rho) \otimes |0\rangle\langle 0|_{\text{f}} \\ &+ p\mathcal{D}_Y(\rho) \otimes |+1\rangle\langle +1|_{\text{f}} \\ &+ (1-p)q\mathcal{D}_Z(\rho) \otimes |-1\rangle\langle -1|_{\text{f}}\,, \end{aligned} \tag{4.5}$$

which outputs on $\mathcal{H}_{\text{s}} \otimes \mathcal{H}_{\text{f}}$ and $\mathcal{D}_\sigma(\cdot)$ is the complete dephasing channel in the $\sigma = \{X, Y, Z\}$ basis, i.e., $\mathcal{D}_\sigma(\rho) = \frac{1}{2}\rho + \frac{1}{2}\sigma\rho\sigma$. In words, this channel describes a flagged dephasing channel,

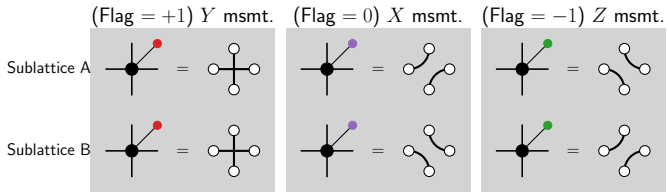
where the local dephasing is done with probabilities $(1-p)(1-q),\ p,\ (1-p)q$ in the X, Y, Z bases respectively.

After the application of this channel, the resulting ensemble is,

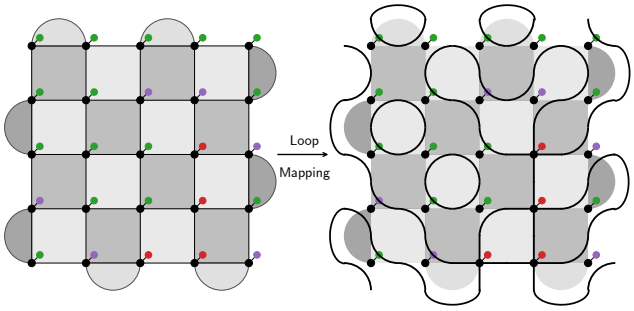
$$\rho_{p,q} := \mathcal{N}_{p,q}^{\otimes n}(\rho_{\text{code}}), \tag{4.6}$$

where ρ_{code} is a surface-code ground state. The state $\rho_{p,q}$ is diagonal in a local product basis; it has no entanglement, but, as we show, it can exhibit nontrivial long-range structure quantified by CMI.

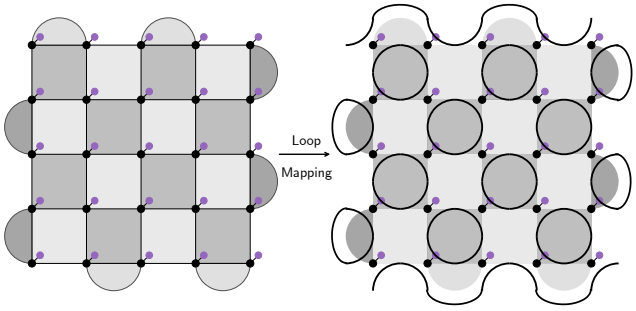
The ensemble $\rho_{p,q}$ can also be interpreted as a result of non-selective single qubit Pauli measurements on the surface code. When such measurements are performed on a subset of the qubits, the remaining unmeasured qubits undergo an entanglement phase transition, which was studied in [165]. Robustness of the stored logical information to single qubit measurements on the surface code was studied in [24, 63, 127].



(a) Local Majorana pairings in the hybrid fermion–qubit tensor for different flag values (measurement bases) on the two sublattices.



(b) Gluing these local tensors via Gaussian contractions yields a configuration of completely packed loops on the lattice.



(c) In the clean memory limit $p = q = 0$, the loop mapping exhibits only plaquette loops, and open strands along the top and bottom boundaries describing the representative logical operators.

Figure 4.3: Mapping from the dephased surface code to loop model.

4.3 Phase diagram of CMI and completely packed loop model

In this section we analyze the CMI in the classical ensemble $\rho_{p,q}$, and relate it to the loop-spanning properties of the completely packed loop model with crossings (CPLC). This allows us to identify the Goldstone phase of the loop model with a critical mixed-state phase with divergent Markov length.

4.3.1 Mapping to loop model

We label lattice sites by $i = 1, \ldots, n$ and denote by $f = (f_1, \ldots, f_n) \in \{-1, 0, +1\}^n$ a configuration of local flags, which specifies the Pauli basis in which each data qubit is dephased. Since the noise channel acts independently on each site, the probability of a given flag configuration factorizes as $p(f) = \prod_i p_{f_i}$, with $p_{f_i} = (1-p)(1-q)$, p, or $(1-p)q$ for $f_i = 0, +1, -1$, respectively. $\rho_{p,q}$ is block diagonal and can be written as

$$\rho_{p,q} = \sum_f p(f)\ |f\rangle\langle f|_{\mathrm{f}} \otimes \sigma_f, \qquad \sigma_f = \mathcal{D}_f(\rho_{\mathrm{code}})\,, \tag{4.7}$$

where $\mathcal{D}_f = \bigotimes_i \mathcal{D}_{f_i}$ denotes complete dephasing of the physical qubits in the Pauli bases specified by f. For fixed f, every physical qubit is dephased (equivalently, nonselectively measured) in a single-qubit Pauli basis $\sigma_i \in \{X_i, Y_i, Z_i\}$. Thus $\sigma_f = \mathcal{D}_f(\rho_{\mathrm{code}})$ is a classical mixture over measurement trajectories or Pauli eigenvalue strings $m_i \in \{\pm 1\}$, and it is maximally mixed on the subspace compatible with the stabilizer constraints that remain after imposing these local measurements. Equivalently, if we denote by $\mathcal{S}_f$ the subgroup of the surface-code stabilizer group that is diagonal in the measured product basis specified by f, then $\sigma_f \propto \prod_{g \in \mathcal{S}_f} \frac{1+g}{2}$, so the classical distribution of outcomes is uniform over assignments satisfying all independent parity checks $\mathcal{S}_f$.

Now we use the Majorana fermionic representation of surface code to show $\mathcal{S}_f$ constraint are given by loop configurations in the loop model. Following the discussion in Sec. 4.2.1, for a given f, the physical leg of each local qubit-fermion tensor is projected onto a σ_i eigenstate, leaving a purely fermionic (Gaussian) stabilizer state on that site. Depending on whether $\sigma_i = X, Y$, or Z (and depending on the sublattice A or B, which are related by a Hadamard), the remaining fermionic state enforces one of a small set of local "rewirings" (pairings) among the four virtual Majoranas; see Fig. 4.3(a).

Contracting neighboring sites with the fixed Gaussian bond projectors glues these local pairing patterns into configurations of completely packed loops on the square lattice [165],

as illustrated in Fig. 4.3(b). Each measurement trajectory with $m_i = \pm 1$ fixes the signs associated with the fermionic pairings at each site, but the overall connectivity pattern remains the same for all trajectories consistent with the flag pattern f.

Each closed loop L appearing in a particular configuration yields a $\mathbb{Z}_2$ constraint on the classical measurement outcomes associated with that loop. The origin of this constraint can be understood by multiplying the local qubit–fermion stabilizers associated with the sites and the bond projectors along the loop, which cancels all internal Majorana operators (each virtual Majorana appears exactly twice), leaving a purely physical Pauli stabilizer operator that survives the local measurements. Importantly, this physical operator is supported only on those sites where the loop contributes an *odd* number of physical Pauli factors: if the loop passes through a site twice, the corresponding Pauli appears twice and cancels ($\sigma_i^2 = \mathbb{I}$), so that site does not appear in the support.

Equivalently, the loop defines a stabilizer element in the subgroup $\mathcal{S}_f$ that is diagonal in the measured product basis, $\prod_{i\in\mathrm{supp}(L)}\sigma_i \in \mathcal{S}_f$, where $\mathrm{supp}(L)$ is the set of sites where the induced physical Pauli appears once and not twice. Evaluating this operator on an outcome string gives the corresponding parity constraint that the measurement outcomes must satisfy. Note that the product of Pauli outcomes around a loop can acquire a sign. In particular,

$$\prod_{i\in\mathrm{supp}(L)} m_{\sigma_i} = s(L) \in \{\pm 1\}, \tag{4.8}$$

where $s(L)$ is the stabilizer phase of the loop operator. A convenient way to express it in the present setting is via the number of Y-type factors that survive on the loop: let $N_Y(L)$ be the number of vertices $i \in \mathrm{supp}(L)$ with $\sigma_v = Y$ (so vertices visited twice by the loop do not contribute). Then the overall phase is $s(L) = (-1)^{N_Y(L)/2}$, since each surviving Y contributes a factor of i when rewriting the loop operator as a product of underlying X- and Z-type stabilizers, and for closed loops $N_Y(L)$ is even. Note, the logical operators are represented as open-strands in the loop configuration, which also imposes a parity constraint on the Pauli outcomes along the strand (see Fig. 4.3c).

The probability of each measurement trajectory $\{m_i\}$ for a particular flag configuration vanishes unless all loop parities (4.8) are satisfied. Moreover, conditioned on a fixed loop configuration, these loop parities generate all surviving constraints, so the allowed outcomes are uniformly distributed over the satisfying constraints. These statements were already derived within the stabilizer formalism; we emphasize the loop interpretation because it makes the connection to the CPLC model explicit.

The CPLC model has a well established phase diagram [160], which is shown in Fig. 4.4. In particular, as a function of p, q, the model hosts two topologically distinct "short loop"

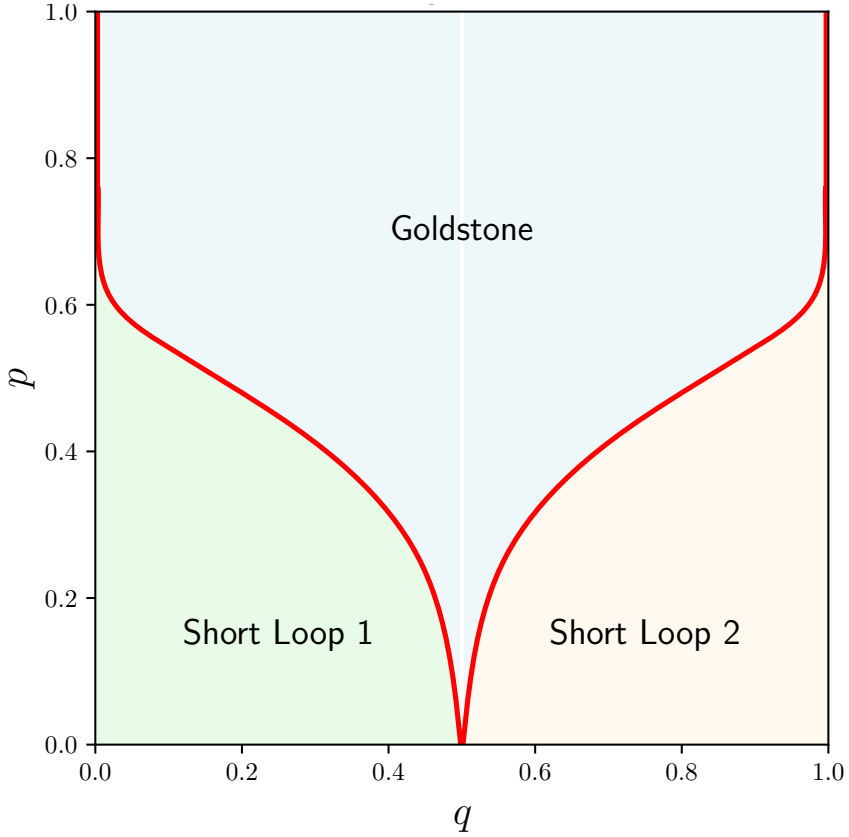


Figure 4.4: Phase diagram for the CPLC model [160].

phases where the typical loops are short, and a "Goldstone" phase where the typical loops span across the system. The phase diagram can be described field theoretically as a $O(n)$ magnet coupled to a $\mathbb{Z}_2$ gauge field in the limit $n \to 1$. The Goldstone phase arises as a Kosterlitz-Thouless type superfluid whose quasi-long range order is destroyed by a proliferation of $\mathbb{Z}_2$ vortices in the short loop phases [160]. Next, we relate the loop spanning properties of the phases to their CMI behavior.

4.3.2 CMI computation

We now fix a tripartition $A{:}B{:}C$ of the physical qubits. Because the flags are classical variables that are statistically independent between sites, and hence between the three regions, the CMI of $\rho_{p,q}$ simplifies considerably. The block-diagonal structure in the flag basis

implies that all von Neumann entropies entering the definition of $I(A{:}C \,|\, B)$ decompose into a flag average of entropies of σ_f, following Eq. 4.7. The classical flag contributions cancel exactly in the CMI, leaving

$$I(A{:}C \,|\, B)_{\rho_{p,q}} = \sum_f p(f)\, I(A{:}C \,|\, B)_{\sigma_f}. \tag{4.9}$$

Thus, the CMI of the full classical ensemble is simply the average, over flag configurations, of the CMI computed for each fixed dephasing pattern f.

Equation (4.9) reduces the analysis of conditional correlations in $\rho_{p,q}$ to the study of the classical distributions σ_f. In the following, we evaluate $I(A{:}C \,|\, B)_{\sigma_f}$ for a fixed flag configuration using the Majorana representation, which maps each σ_f to a constrained loop model determined by the spatial pattern of local Pauli dephasing, as described in the earlier subsection.

For a particular non-overlapping tripartition $A{:}B{:}C$ of the qubits on the surface code patch, the flag configurations can be organized into two cases, one where there are no loops (constraints) that simultaneously touch A and C, and the other where there are loops that do, as can be seen in the following schematic in Fig. 4.5a.

First note that, only those flag patterns whose induced loop configuration contains at least one loop that touches both A and C can contribute to the CMI. If every loop that intersects A is entirely contained in $A \cup B$, and every loop that intersects C is entirely contained in $B \cup C$, then $I(A : C \mid B) = 0$ for that configuration. In contrast, a loop that traverses from A to C through B generates a $\mathbb{Z}_2$ parity constraint relating measurement outcomes in A and C that is not removed by conditioning on B. The CMI (in log base 2) for a fixed flag pattern is therefore equal to the number of *independent* A–C parity constraints that survive after conditioning on B. Consequently, the CMI is exactly governed by the loop statistics of the CPLC model.

For a fixed flag pattern f, write all $\mathbb{Z}_2$ parity constraints of σ_f as a binary matrix $C_{ABC}(f)$ (rows = constraints, columns = qubits). Conditioning on B amounts to fixing the B bits, so for the purpose of counting independent constraints we may delete the B columns. Let $C_{AC}(f)$ be the resulting matrix, and let $C_A(f)$ and $C_C(f)$ denote its restrictions to the A and C columns, respectively. Then the number of *irreducible* constraints that have support on both A and C (i.e. not generated by A-only and C-only constraints) is

$$\begin{aligned} I(A{:}C \mid B)_{\sigma_f} = {} & \operatorname{rank}_{\mathbb{F}_2} C_A(f) + \operatorname{rank}_{\mathbb{F}_2} C_C(f) \\ & - \operatorname{rank}_{\mathbb{F}_2} C_{AC}(f). \end{aligned} \tag{4.10}$$

where the ranks are obtained by Gaussian elimination over $\mathbb{F}_2$.

We begin with the *pairwise* geometry in Fig. 4.5b, where A and C each consist of a single qubit and B is the rest. In this case, the only way to obtain a constraint between A and C after conditioning on B is for the loop configuration to contain *two distinct loops* that connect A to C (see Fig. 4.5c right). There are other possibilities—A and C may not be connected at all, or they may be connected by a single loop that visits each site once or twice (see Fig. 4.5c left)—but in either case no parity constraint with net support on both A and C survives after conditioning on B, and therefore these configurations do not contribute to $I(A{:}C \mid B)_{\sigma_f}$.

It is useful to translate this into standard loop-model language. The relevant observable is the *four-leg* ("watermelon") correlator $G_4(r)$ [160], defined as the probability that two marked sites separated by distance r are connected by a total of four strands. Importantly, the event "four strands connect A and C" has two distinct topological components, depending on how the strands are paired internally at the marked sites. In one component, the four strands belong to a *single* loop that visits both A and C twice, so that A and C are connected through one large loop. In the other component, the four strands split into *two distinct* loops, each of which visit A and C once. We therefore decompose $G_4(r) = G_4^{(1)}(r) + G_4^{(2)}(r)$, where $G_4^{(1)}$ is the single-loop contribution and $G_4^{(2)}$ is the two-loop contribution, which are shown in Fig. 4.5c.

In the pairwise geometry, only the two-loop component produces an A–C parity constraint that survives conditioning on B, and hence

$$I(A{:}C \mid B)_{\text{pairwise}} = G_4^{(2)}(r), \tag{4.11}$$

i.e., the pairwise CMI is exactly the flag-averaged probability that A and C are connected by two distinct loops. As argued in Ref. [160], the two-loop component $G_4^{(2)}(r)$ has the same long-distance scaling as the full four-leg correlator $G_4(r)$, so the pairwise CMI inherits the same asymptotic scaling behavior. Hence $I(A:C \mid B) \asymp G_4(r)$, with scaling

$$I(A:C \mid B)_{\text{pairwise}} \sim \begin{cases} r^{-2x_4}, & \text{transition}, \\ e^{-r/\xi_{\text{loop}}}, & \text{short loop}, \\ [\log(r/l)]^{-12}, & \text{Goldstone}, \end{cases} \tag{4.12}$$

where x_4 is reported numerically in Ref. [160].

For the *local* and *global* geometries, the CMI is most conveniently evaluated by directly implementing Eq. (4.10) for each sampled flag pattern: we determine the loop-parity constraint space and compute the number of independent A–C constraints that remain after conditioning on B using the stabilizer formalism (Gaussian elimination over $\mathbb{F}_2$). Figure 4.6

shows the resulting behavior at fixed $q = 0.1$ as a function of p. For the pairwise geometry, this numerical procedure provides a direct check of the loop-correlator prediction Eq. (4.11) and its implied scaling Eq. (4.12).

We can bound the local CMI by the pairwise CMI and statistics of loop lengths in the Goldstone phase. Consider a local geometry (Fig. 4.5b) where A is a single qubit, and the region C is decomposed into a single qubit C_1 at a distance r from A and the rest $C_2 = C \setminus C_1$. By the chain rule of CMI and strong subadditivity, we have

$$I(A : C|B)_{\text{local}} = I(A : C_1C_2|B) \geq I(A : C_1|C_2B), \tag{4.13}$$

where the RHS is exactly the pairwise CMI. Thus, in the Goldstone phase, we must have at least as slow a decay as the pairwise CMI i.e. $\log(r/l)^{-12}$.

In the Goldstone phase, Ref. [160] showed that the non-returning probability $Q(l)$ for a loop through a fixed link after l steps (their Eq. (52)) decays as $Q(l) \sim \text{const}/\ln(l/l_0)$. Hence the probability $\pi(r)$ that such a loop makes an excursion to distance at least r is bounded by $\pi(r) \leq Q(r) \lesssim \text{const}/\ln(r/r_0)$. For the local geometry, a necessary condition for nonzero $I(A{:}C|B)_{\text{local}}$ is that some loop incident on A reaches C across the buffer, which implies that at least one of the $O(1)$ links in ∂A has an excursion of size $\geq r$. By a union bound this event has probability $\lesssim \text{const}/\ln(r/r_0)$ and the $I(A{:}C|B)_{\text{local}}$ is upper-bounded by this probability. Putting these bounds together, we infer

$$I(A{:}C|B)_{\text{local}}(r) \;\sim\; \frac{1}{[\log(r/l)]^{\alpha}}, \tag{4.14}$$

where $1 \leq \alpha \leq 12$, up to nonuniversal prefactors and subleading corrections. In particular, the decay is slower than any power law and the Markov length extracted from the local geometry diverges. Numerically, the local CMI is shown to follow a similar transition as the pairwise CMI, as presented in Fig. 4.6b.

In the *global* geometry, the CMI exhibits a clear transition (for $q = 0.1$) at $p_c \simeq 0.571$ with a finite-size scaling collapse consistent with $\nu \simeq 2.79$ [Fig. 4.6a]. For $p < p_c$ the global CMI decays exponentially with L, while for $p > p_c$ it grows logarithmically with L, as expected in the short-loop and Goldstone regimes, respectively. In the *local* and *pairwise* geometries, fixing $d_{AC} = L/2$ we find exponential decay with L for $p < p_c$, while for $p \geq p_c$ the decay becomes sub-exponential [Fig. 4.6b], consistent with the interpretation of pairwise CMI as the $G_4^{(2)}$ component of the four-leg correlator (and hence with the asymptotic scaling inherited from G_4). Along the $p = 0$ line, the transition at $q = 0.5$ is bond percolation with $\nu = 4/3$ [160]; away from $p = 0$ the transition leaves the percolation universality class, consistent with the exponent extracted from the data in Fig. 4.6.

From Fig. 4.6a we see that the global CMI exhibits the same qualitative behavior across the phase diagram as the loop-model *spanning number*. Define $N_{\text{span}}(A \leftrightarrow C)$ to be the number of loop strands that connect A to C through B in a given loop configuration (and let $\langle\cdot\rangle_f$ denote the average over flag patterns, equivalently over loop configurations). The spanning number is a natural diagnostic of the short-loop versus Goldstone regimes: in the short-loop phase it is exponentially suppressed with separation, while in the Goldstone phase it grows logarithmically, reflecting the proliferation of long wandering strands [160].

N_{span} provides an upper bound to the global CMI. Each independent A–C parity constraint that survives conditioning on B must be supported on at least two A–C spanning strands, which implies the upper bound $I(A{:}C \mid B)_{\text{global}} \leq \frac{1}{2}\left\langle N_{\text{span}}(A \leftrightarrow C)\right\rangle_f$. In the short-loop phase, where both quantities decay exponentially in the separation r, this immediately implies that the Markov length ξ_{M} extracted from $I(A{:}C \mid B)_{\text{global}} \sim e^{-r/\xi_{\text{M}}}$ is upper bounded by the loop length scale ξ_{loop} extracted from $\langle N_{\text{span}}(A \leftrightarrow C)\rangle_f \sim e^{-r/\xi_{\text{loop}}}$, i.e. $\xi_{\text{M}} \leq \xi_{\text{loop}}$.

Remarkably, our numerics indicate that this bound is saturated throughout the short-loop phase: the decay lengths extracted from the global CMI and from the spanning number coincide within numerical accuracy, $\xi_{\text{M}} = \xi_{\text{loop}}$ for $p < p_c$, as shown in Fig. 4.7. This provides direct evidence that the Markov length of the classical ensemble is set by the same loop correlation length controlling strand-spanning in the CPLC model.

Combining (i) the numerical scaling of the global CMI across the transition in Fig. 4.6a with (ii) the numerical observation that these two decay lengths coincide throughout the short-loop regime, $\xi_{\text{M}} = \xi_{\text{loop}}$ (Fig. 4.7), we arrive at the scaling forms

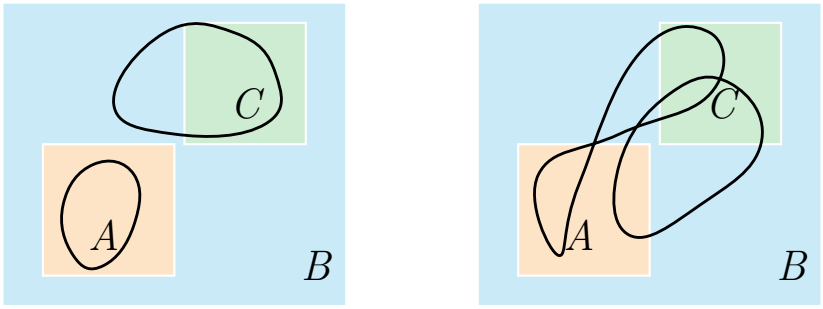
$$I(A{:}C \mid B)_{\text{global}} \sim \begin{cases} \text{constant}, & \text{transition}, \\ e^{-r/\xi_{\text{loop}}}, & \text{short loop}, \\ \log(r/l), & \text{Goldstone}. \end{cases} \tag{4.15}$$

Our numerical results confirm that the CMI behavior across all geometries follows the CPLC phase diagram in Fig. 4.4. In the short loop phase, the CMI in all geometries are exponentially decaying, implying the existence of a finite Markov length, while at the critical point and the Goldstone phase the Markov length diverges. We quantitatively do not find the Markov length obtained from the local, pairwise, and global geometries to be equal; however, that may arise from the fact that the different geometries could come with distinct poly(r) factors as well.

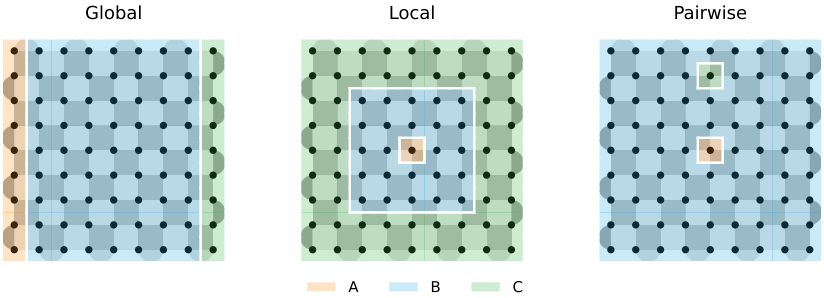
In summary, for the classical ensemble $\rho_{p,q}$, we find that the CMI behavior mirrors the

loop spanning behavior of the underlying loop model, which exhibits stable critical and short-loop phases and a transition between them. [1]

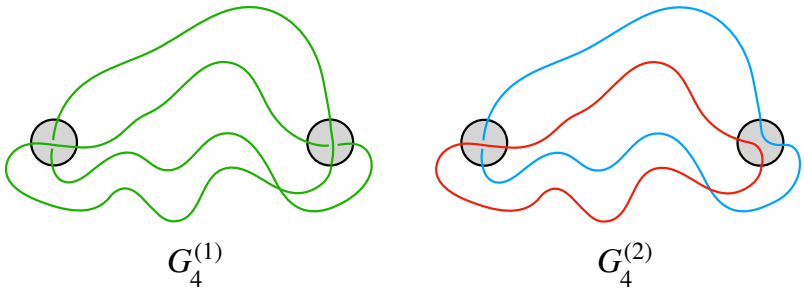
[1] We have also studied the CMI and other informational quantities along the line $q = 0.0$, as a function of p, and find an apparent crossover at $p^* < 1.0$. As reported in [160], the localization length diverges near the $p = 1.0$ fixed point, which leads to strong finite size constraints in the numerical analysis. We report the results in the Appendix C.2, and perform finite size scaling to show that this crossing is not a real transition.



(a) Loop configurations that do not (left) and do (right) contribute to the CMI $I(A:C|B)$.

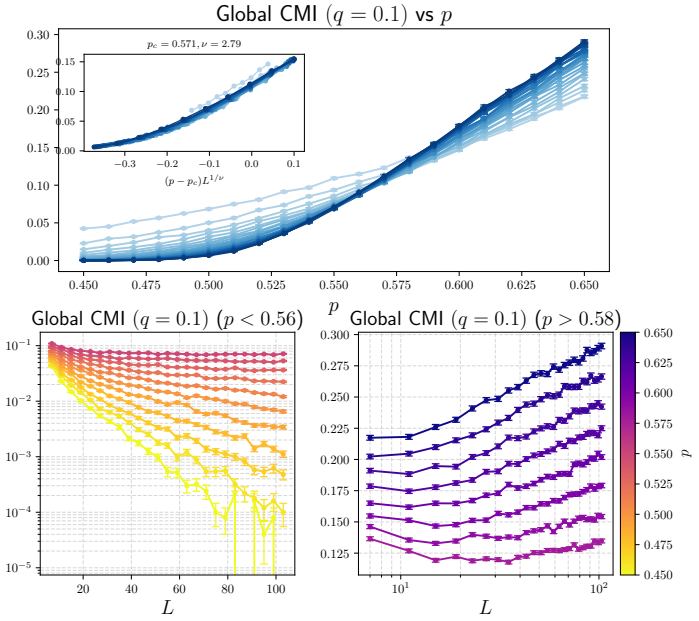


(b) Surface code geometries of tripartition for which the CMI is computed for the classical ensemble: “global", “local", and “pairwise".

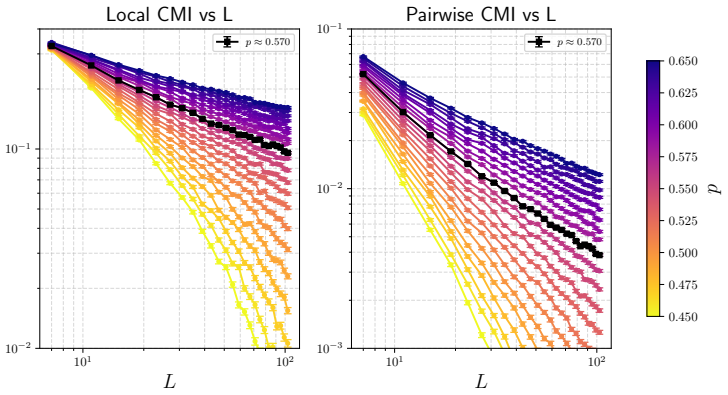


(c) The two topologically distinct diagrams for the “pairwise" strand counting. Only $G_4^{(2)}$ contributes to the CMI.

Figure 4.5: Loop configurations and geometries for the CMI computation.



(a) The global CMI as a function of p for a fixed $q = 0.1$, and for different sizes of the surface code patch of $L \times L$, with $L = 21 - 161$. There is a clear transition at around $p_c \approx 0.571$, with the critical exponent $\nu \approx 2.79$. The traces are averaged over 5×10^4 samples.



(b) The CMI of the classical ensemble for local (left) and pairwise (right) geometries, with $d_{AC} = L/2$ is plotted as a function of L on the log-log scale. Both decay exponentially with L for $p < p_c$, and sub-exponentially with L for $p > p_c$.

Figure 4.6: Numerical results for the CMI for different geometries for a fixed $q = 0.1$, with varying p and sizes of the surface code patch L, which diagnose the short loop to Goldstone phase transition in the CPLC model.

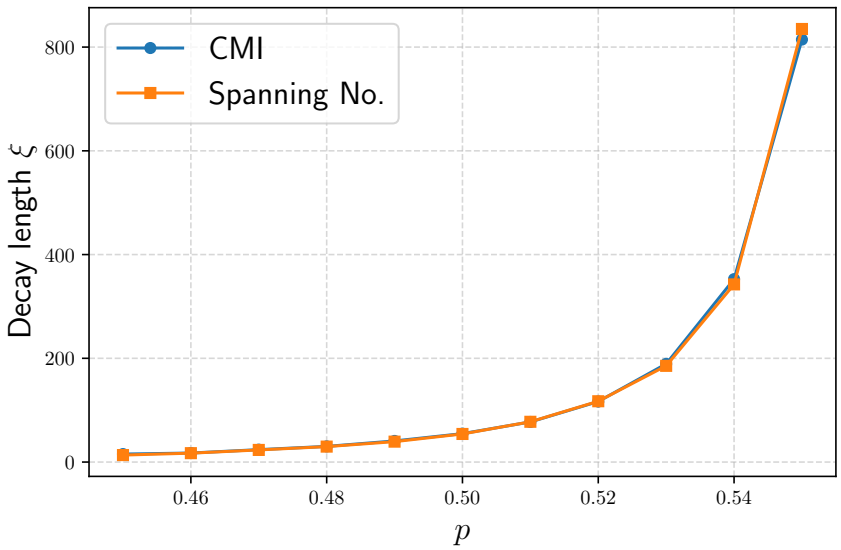


Figure 4.7: In the short-loop phase $p < p_c$, the decay length extracted from the global CMI (the Markov length ξ_{M}) coincides within numerical accuracy with the loop length scale ξ_{loop} extracted from the spanning number.

4.4 Critical phase as a memory

In this section we ask what the extended critical mixed-state phase implies for fault tolerance. From error correction perspective there are two distinct questions: (i) whether logical information is still present after the noise, and (ii) if it is present, whether it can be recovered by a quasi-local decoder or only by intrinsically global processing. To address (i), we use the mutual information between a reference and the noisy code block, $I(R{:}Q)$, which plays the role of an input output information measure for the noise channel. To address (ii), we refine this diagnostic by introducing a "punctured" version, obtained by discarding a bulk region of the code before computing the mutual information. Comparing the full and punctured quantities provides a direct operational probe of *local recoverability*: if decoding can be implemented by a quasi-local circuit, then removing a sufficiently small interior region should not affect the information accessible from the remaining degrees of freedom, whereas a puncture-induced drop signals that any successful decoder must be nonlocal. We now formalize these notions in the classical setting relevant for storing a logical bit.

4.4.1 Memory and the recovery

Complete X-dephasing of the surface-code ground state produces a classical ensemble stabilized by the X-type plaquette checks. In this ensemble the only remaining logical degree of freedom is the logical operator $\bar{X}$, so the code stores a single classical bit. We couple this bit to a classical reference register R by enforcing $X_r\bar{X} = 1$, without imposing any Z-type constraint. The corresponding reference–code output state is defined by

$$\rho_{RQ} := (\mathrm{id}_R \otimes \mathcal{D}_X)\left[\left(\frac{1 + X_r\bar{X}}{2}\right)\left(\mathbb{I}_R \otimes \rho_{\mathrm{code}}\right)\right], \tag{4.16}$$

where $\mathcal{D}_X$ completely dephases the code in the X basis and ρ_{code} is chosen to have fully mixed logical space. The marginal ρ_Q is the classical surface-code ensemble storing the logical X bit.

The corresponding noisy reference–code output state is defined, for parameters (p, q), by

$$\rho_{RQ}(p, q) := (\mathrm{id}_R \otimes \mathcal{N}_{p,q}^{\otimes n})\left[\left(\frac{1 + X_r\bar{X}}{2}\right)\left(\mathbb{I}_R \otimes \rho_{\mathrm{code}}\right)\right], \tag{4.17}$$

The Pauli dephasing channel $\mathcal{N}_{p,q}^{\otimes n}$ acts on Q to produce the classical ensemble $\rho_{p,q}$ defined in Eq. (4.6). Since $\mathcal{N}_{p,q}^{\otimes n}$ outputs both the data qubits and a classical flag register

F indicating the local Pauli basis, we will henceforth include the flags in the output system and write $Q \equiv Q_{\text{data}} \otimes F$. Equivalently, the noise channel can be viewed as performing nonselective single-qubit Pauli measurements, with the flags revealing when a Y- or Z-basis measurement occurred instead of an X-basis measurement.

We say that the classical logical bit is recoverable if there exists a recovery channel that can map the noisy state to the clean state: $\mathcal{R}$ on Q such that

$$(\mathrm{id}_R \otimes \mathcal{R})\big(\rho_{RQ}(p,q)\big) \;\to\; \rho_{RQ}(0,0). \tag{4.18}$$

A finer question is whether recovery can be implemented by a *local* decoder. To make this precise, we consider recovery maps $\mathcal{R} : Q \to Q$ acting on the physical degrees of freedom equipped with the lattice metric. We say that $\mathcal{R}$ is *local* if it can be implemented as a channel circuit of bounded depth, where each layer is a tensor product of CPTP maps supported on regions of bounded diameter (equivalently, contained in the neighborhood of some site). We will be interested both in *strictly local* recovery, where the circuit depth and gate diameter are $O(1)$, and in *quasi-local* recovery, where they are allowed to scale polylogarithmically with the linear system size L, i.e., $O\big((\log L)^k\big)$ for some fixed k. This locality constraint directly probes the circuit complexity of decoding and is therefore a sharper notion than global recoverability.

Our notion of local recovery map is closely related to the *local reversibility* framework developed in Ref. [193]. There, one considers a setting in which a state is generated by evolution under a local Lindbladian. Ref. [193] shows that if the Markov length (defined via exponential decay of the CMI for annular partitions) remains finite along the Lindbladian trajectory, then the evolution is reversible by another quasi-local Lindbladian, with the locality scale controlled by the Markov length along the path. Equivalently, finite Markov length along the noise evolution provides a *sufficient* condition for the existence of a quasi-local recovery/decoder.

4.4.2 Diagnostics of global recoverability

We first focus on *global* recoverability of the stored classical bit. The natural diagnostic is the input–output mutual information[2] $I(R{:}Q)$ [202, 203, 95, 100, 128], which is a standard decoder-independent measure of how much information about the logical input can still be inferred from the noisy code and, equivalently, whether the output retains correlations

[2]This is the coherent information $I_c(R\rangle Q) \coloneqq S(Q) - S(RQ)$ plus a constant equal to the number of logical qubits; therefore, we use coherent and mutual information interchangeably in this context.

with the logical reference. If $I(R{:}Q)$ saturates to its maximal value of one bit, the logical information is perfectly preserved and hence globally recoverable; if $I(R{:}Q)$ vanishes in the thermodynamic limit of the system, then the logical information has been lost to the environment and cannot be recovered by any decoder. Concretely, after the flagged Pauli-measurement noise, $I(R{:}Q)$ quantifies how much information about the logical input remains accessible from the *entire* output, assuming an ideal decoder with unrestricted access to all degrees of freedom:

$$I(R{:}Q) = S(R) + S(Q) - S(RQ), \tag{4.19}$$

with logarithms in base 2. Here we take $Q \equiv Q_{\text{data}} \otimes F$, i.e. the output includes both the data qubits and the classical flags.

The flagged noise is a classical mixture of nonselective single-qubit Pauli measurements, so each flag sample specifies a full measurement trajectory f (which basis $\sigma_i \in \{X, Y, Z\}$ was applied at every site). For fixed f, the trajectory induces an effective logical nonselective measurement $\bar{\sigma}(f) \in \{\bar{X}, \bar{Y}, \bar{Z}\}$. Since the reference is correlated only with the logical $\bar{X}$ bit, the trajectory mutual information is binary: $I(R{:}Q)_f = 1$ if $\bar{\sigma}(f) = \bar{X}$ and $I(R{:}Q)_f = 0$ otherwise. Averaging over trajectories gives

$$I(R{:}Q) = \sum_f p(f)\, I(R{:}Q)_f = \Pr\big(\bar{\sigma}(f) = \bar{X}\big). \tag{4.20}$$

To evaluate the right-hand side, we use the loop-model mapping. A fully measured bulk induces a pairing of the four dangling boundary Majorana modes of the surface-code patch. There are three topologically distinct pairings, and they correspond to effective logical measurement bases Y, X, and Z, as shown in Fig. 4.8. Only the X-type pairing preserves the encoded classical $\bar{X}$ bit. Therefore, Eq. (4.20) can be computed by sampling trajectories and recording which boundary pairing sector they induce; $I(R{:}Q)$ is the frequency of the X-type pairing.

The resulting $I(R{:}Q)$ obtained from Monte Carlo sampling is shown in Fig. 4.9. It exhibits three regimes, $I(R{:}Q) \to 1$, $I(R{:}Q) \to 0$, and $I(R{:}Q) \to 1/3$; respectively, these coincide with the two short-loop phases and the intervening Goldstone phase of the loop model discussed in Sec. 4.3. In one short-loop phase the boundary sector is pinned to the X-type pairing, so typical trajectories measure $\bar{X}$ and $I(R{:}Q) \to 1$. In the other short-loop phase the boundary sector is pinned to the Z-type pairing, so typical trajectories measure $\bar{Z}$ and the classical $\bar{X}$ information is erased, giving $I(R{:}Q) \to 0$. In the Goldstone phase the boundary mode delocalizes into the bulk and the boundary sector is not pinned; in

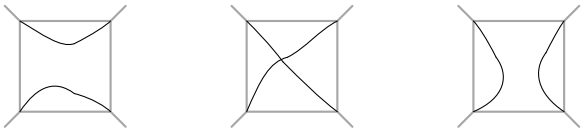


Figure 4.8: Three topologically distinct pairings of the four boundary Majorana modes induced by a trajectory, corresponding to effective logical measurement bases X, Y, and Z (left to right). Only the X-type pairing preserves the encoded classical $\bar{X}$ bit.

the large-distance limit the three topological pairing sectors occur with equal probability, which yields $I(R{:}Q) \to 1/3$ via Eq. (4.20).

It is useful to translate the mutual information $I(R{:}Q)$ into an operational decoding statement. Since the encoded information is a single classical bit, the performance of any decoder can be characterized by its *success probability* p_{succ}, defined as the probability that the decoder correctly outputs the logical value stored in the reference. In the usual error-correction setting, one encounters two limiting behaviors in the thermodynamic limit: if the logical information is perfectly preserved, then $I(R{:}Q) \to 1$ and an optimal decoder succeeds with probability $p_{\text{succ}} \to 1$; if the logical information is completely lost, then $I(R{:}Q) \to 0$ and the best any decoder can do is random guessing, which succeeds with probability $p_{\text{succ}} = 1/2$ for a classical bit. In both cases, increasing the code distance drives the success probability exponentially close to these limiting values.

The Goldstone phase realizes a qualitatively different regime. Here the logical information is neither fully preserved nor completely erased: instead, the flagged measurement trajectories induce the three logical sectors $\{\bar{X}, \bar{Y}, \bar{Z}\}$ with equal probability. With probability $1/3$ the trajectory lies in the $\bar{X}$ sector, in which case a global decoder has full access to the logical bit and succeeds with unit probability. With the remaining probability $2/3$, the trajectory lies in a $\bar{Y}$ or $\bar{Z}$ sector, for which the output is independent of the encoded $\bar{X}$ bit and the optimal decoder can do no better than random guessing, succeeding with probability $1/2$. As a result, the optimal decoding success probability in the Goldstone phase is $p^{\star}_{\text{succ}} = \frac{1}{3} \cdot 1 + \frac{2}{3} \cdot \frac{1}{2} = \frac{2}{3}$. Thus a decoder can only recover with a fixed probability strictly between the two conventional extremes 1 and $1/2$. In the next section we ask whether any local decoder can achieve an asymptotic success probability larger than $1/2$. We answer this question in the negative, and show that no quasi-local decoder can outperform random guessing, despite the existence of a global decoder with success probability $2/3$.

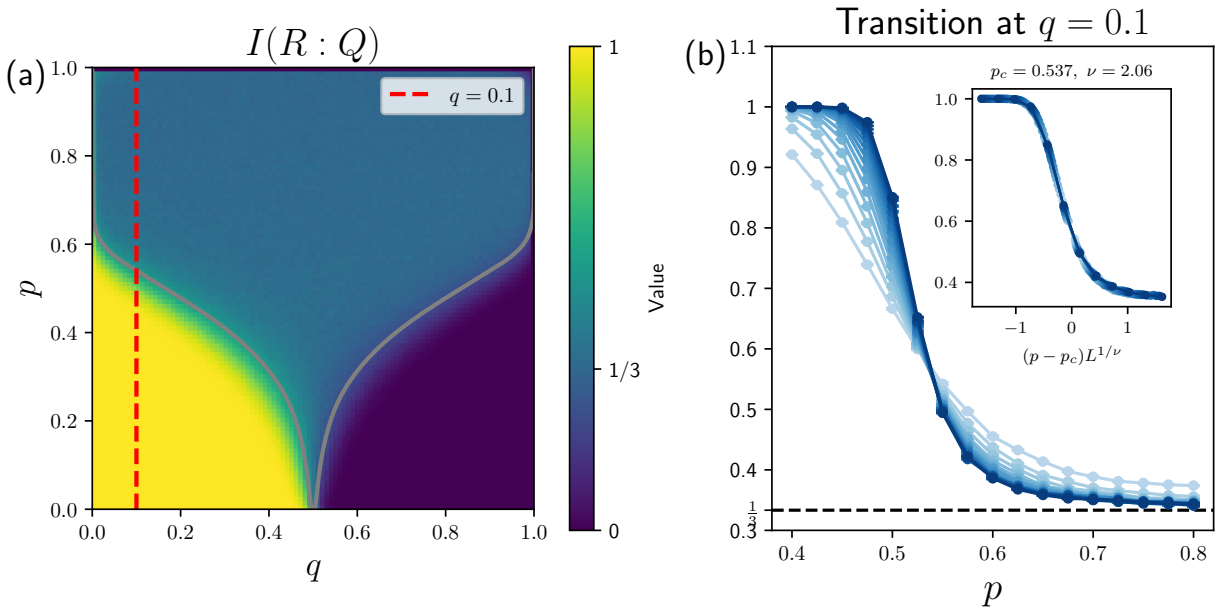


Figure 4.9: (a) Heat map of the input–output mutual information $I(R{:}Q)$ in the (p, q) plane for a surface code patch of linear size $L = 101$, computed numerically by sampling flagged measurement trajectories and identifying the induced boundary pairing sector as in Fig. 4.8. The three plateaus $I(R{:}Q) \to 1$, $1/3$, and 0 correspond to the short-loop–X, Goldstone, and short-loop–Z regimes of the loop model with crossings. Note, the grey line is a guide to the eye and not a phase boundary, rather a contour line at $I(R:Q) = 0.5$.(b) One-dimensional cut at fixed $q = 0.1$, showing $I(R{:}Q)$ versus p for several system sizes $L = 21 - 161$ with 10^4 samples, together with a finite-size scaling collapse. The crossing and collapse locate the transition at p_c and give a critical exponent $\nu \simeq 2$.

4.4.3 Diagnostics of local recoverability

We now address a finer question than global recoverability: can the logical information be recovered by a decoder of bounded circuit complexity, i.e., by a local or quasi-local quantum circuit as previously defined? In general, there is no known necessary and sufficient information-theoretic criterion for quasi-local recoverability. Here we introduce a *necessary* condition for quasi-local recovery, and we will use its contrapositive to certify an intrinsically nonlocal recovery regime in our setting.

The key intuition is that a quasi-local recovery map cannot have arbitrarily long-distance causal influence. In particular, if one removes ("punctures") a region Q_p be-

fore decoding, then a quasi-local decoder can only propagate the effect of this removal to a slightly enlarged neighborhood of the puncture (its causal cone). If this enlarged neighborhood is sufficiently smaller than the code distance (specifically, it cannot support any nontrivial logical operator), then puncturing cannot change the asymptotic reference–output correlations that are accessible from the complement. This motivates the following diagnostic. For each lattice size L, choose a region $Q_p(L) \subseteq Q$ (a "puncture") and let $Q_{p^c}(L) := Q \setminus Q_p(L)$ denote its complement. Given any state ρ_{RQ}, define the punctured mutual information by

$$I(R{:}Q_{p^c}) := S(R) + S(Q_{p^c}) - S(RQ_{p^c}), \tag{4.21}$$

where RQ_{p^c} is the reduced state obtained from ρ_{RQ} by tracing out Q_p. The puncturing test compares the mutual information between R and the full register Q to the mutual information retained after discarding Q_p. The following theorem shows that if recovery can be implemented quasi-locally, then there exist punctures $Q_p(L)$ which grow with system size but nevertheless do not affect the asymptotic reference–output correlations, so that the difference between $I(R{:}Q)$ and $I(R{:}Q_{p^c}(L))$ vanishes in the large-L limit. Before stating the theorem, we define *topological stabilizer codes* as families of stabilizer codes on a D-dimensional lattice with geometrically local stabilizers (bounded support diameter) and distance $d(L) \to \infty$ as $L \to \infty$ [32].

Theorem 4.1 (Punctured coherent/mutual information)**.** *Let $\{\rho_{RQ}^{(L)}\}_L$ be a family of reference–code states on $R \otimes Q$, where $Q = Q(L)$ is the physical register of a family of topological stabilizer codes with distance $d = d(L)$. Let $\rho_{RQ,\mathrm{noisy}}^{(L)}$ be the corresponding noisy reference–output state and let $\rho_{RQ,\mathrm{ideal}}^{(L)}$ be an ideal target state satisfying all stabilizer checks on Q.*

If (i) the noisy and ideal families satisfy the conditions [3]

$$\left| I(R{:}Q)_{\rho_{RQ,\mathrm{noisy}}^{(L)}} - I(R{:}Q)_{\rho_{RQ,\mathrm{ideal}}^{(L)}} \right| \xrightarrow[L\to\infty]{} 0, \tag{4.22}$$

$$\left\| (\mathrm{id}_R \otimes \mathcal{R}_L)\big(\rho_{RQ,\mathrm{noisy}}^{(L)}\big) - \rho_{RQ,\mathrm{ideal}}^{(L)} \right\|_1 \xrightarrow[L\to\infty]{} 0, \tag{4.23}$$

for some family of recovery channels $\{\mathcal{R}_L\}_L$ implementable by a quasi-local circuit, and (ii) there exist constants $c > 0$ and $\alpha > 0$ such that $d(L) \geq cL^\alpha$ for all sufficiently large L, then

[3] If the noisy state is obtained by applying a noise channel acting only on Q to the ideal state, then (4.23) implies (4.22). We state (4.22) separately to allow more general settings in which only part of the reference–output correlations are recoverable, so that the "ideal" state represents the best performance achievable by any decoder. This is particularly relevant in the Goldstone phase, where the output retains only partial logical information.

there exists a family of punctures $\{Q_p(L) \subseteq Q\}_L$ *whose size* $|Q_p(L)|$ *grows polynomially in* L *such that, writing* $Q_{p^c}(L) := Q \setminus Q_p(L)$*,*

$$\left| I(R{:}Q_{p^c})_{\rho^{(L)}_{RQ,\text{noisy}}} - I(R{:}Q)_{\rho^{(L)}_{RQ,\text{noisy}}} \right| \xrightarrow[L\to\infty]{} 0. \tag{4.24}$$

Proof sketch. A region M is correctible iff there are no non-trivial logical Pauli supported solely on M. In particular, any M that cannot support a nontrivial logical operator (e.g. any region of weight $< d$) is correctable by definition of the distance. On a D-dimensional lattice, a hypercubic region of linear size ℓ contains $O(\ell^D)$ qubits, hence it is correctable whenever $\ell < d^{1/D}$. Using $d(L) \geq cL^\alpha$, we may therefore choose punctures $Q_p(L)$ of linear size $\text{diam}(Q_p(L)) = \gamma\, L^{\alpha/D}$ for some constant $0 < \gamma < 1$. Quasi-locality of $\mathcal{R}_L$ implies a bounded light-cone: the effect of removing $Q_p(L)$ can only influence the recovered state within a slightly enlarged neighborhood of $Q_p(L)$ (with at most polylogarithmic overhead in linear size). For γ chosen sufficiently small, this enlarged neighborhood still has size below the distance scale and is therefore correctable. Hence puncturing can only affect the recovered state inside a correctable region and cannot change the recovered state on its complement. Therefore, the reference correlations are asymptotically unchanged, giving (4.24). See Appendix C.1 for details. □

The puncture scaling suggested by the proof sketch should be viewed as a general baseline that follows from locality and polynomially growing distance, without invoking any additional structure of the code family. For the surface code, however, one can choose substantially larger punctures: the mutual-information invariance holds for geometric punctures whose linear size is any fixed fraction of the system size. This improvement reflects two standard features of the surface code: correctable regions can be taken with linear size proportional to the distance, and the distance scales linearly with the lattice size. Together, these observations lead to the following corollary.

Corollary 4.2 (Surface-code puncture invariance)**.** *In the setting of Theorem 4.1, assume the underlying code family is a surface code on a two-dimensional lattice of linear size* L*. Fix any constant* $0 < \gamma < 1$*, and let* $Q_p(L) \subseteq Q$ *be the puncture defined by a square sublattice (a square "hole") of linear size*

$$\text{diam}(Q_p(L)) = \gamma L, \qquad 0 < \gamma < 1. \tag{4.25}$$

Writing $Q_{p^c}(L) := Q \setminus Q_p(L)$*, the punctured mutual information satisfies* $I(R{:}Q_{p^c})_{\rho^{(L)}_{\text{noisy}}} \xrightarrow[L\to\infty]{} I(R{:}Q)_{\rho^{(L)}_{\text{noisy}}}$.

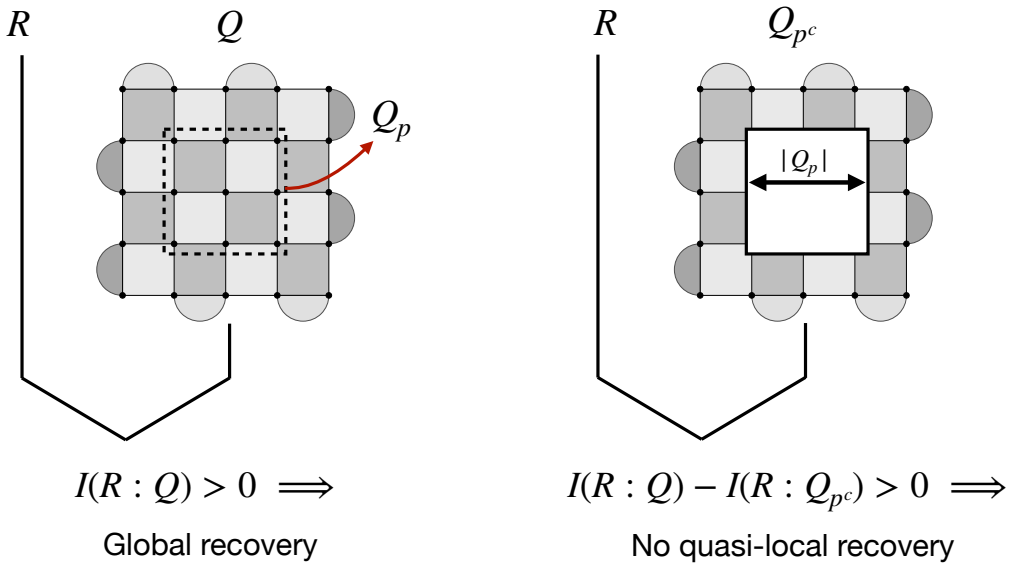


Figure 4.10: Punctured probes of input–output mutual information. If decoding is quasi-local, discarding a finite fraction of system size does not change the accessible information about the logical input. A persistent drop under puncturing therefore witnesses intrinsically nonlocal recovery.

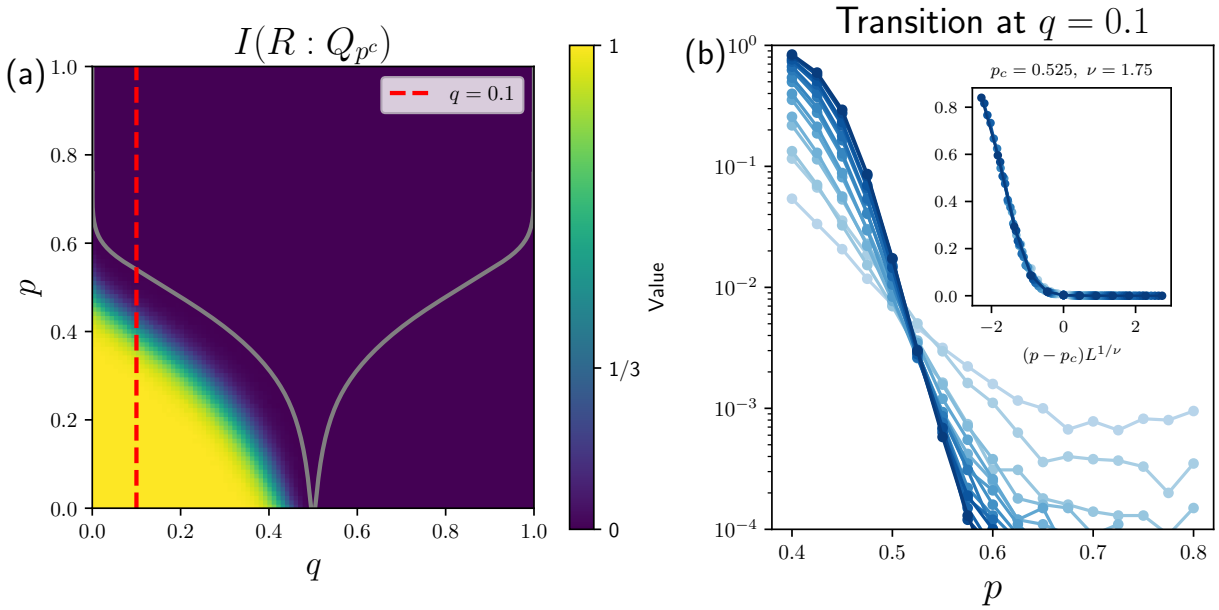


Figure 4.11: Punctured input–output mutual information $I(R{:}Q_{p^c})$ computed using the protocol of Fig. 4.10, with a bulk square puncture of linear size $L/2$. The $I(R{:}Q) \to 1$ and $I(R{:}Q) \to 0$ short loop regimes are stable under puncturing, while the $I(R{:}Q) \to 1/3$ Goldstone regime collapse to $I(R{:}Q_{p^c}) \to 0$. The small mismatch between the apparent transition points for punctured and unpunctured curves is consistent with finite-size drift. The plot parameters are the same as Fig. 4.9. Note that the purple subregion below the transition arises from the fact that the heatmap is plotted for a fixed system size; finite size scaling reveals that the $I(R : Q_{p^c})$ goes to 1 in the $L \to \infty$ limit below the transition, as shown in (b).

No local recovery in the Goldstone phase

We now apply Corollary 4.2 as a diagnostic for nonlocality of decoding in our classical surface-code ensemble. In the numerics, the puncture Q_p is a bulk square of linear size $L/2$, so that a fixed finite fraction of the bulk is discarded. We then compute $I(R{:}Q_{p^c})$ by sampling flagged trajectories exactly as for $I(R{:}Q)$, but tracing out all degrees of freedom (data and flags) supported in Q_p. Figure 4.11 summarizes the results.

The puncturing test cleanly separates the short-loop regimes from the Goldstone phase. In both short-loop phases, puncturing does not change the asymptotic input–output mutual information: in the $\bar{X}$-pinned short-loop phase we have $I(R{:}Q) \to 1$ and also $I(R{:}Q_{p^c}) \to 1$, while in the $\bar{Z}$-pinned short-loop phase we have $I(R{:}Q) \to 0$ and likewise $I(R{:}Q_{p^c}) \to 0$. In particular, in the $\bar{X}$-pinned regime, discarding half of the bulk does not affect the recoverability of the bit. This stability admits a simple loop-model interpretation: viable $\bar{X}$ representatives are confined within distance $\sim \xi_{\text{loop}}$ of the top and bottom boundaries, so the probability that all $\bar{X}$ representatives intersect a bulk puncture that is a macroscopic distance from the boundary is exponentially small in the boundary-to-puncture separation. Equivalently, the event that puncturing changes the decoder-relevant logical sector is controlled by the same exponentially decaying strand-spanning probabilities that define the short-loop regime.

The Goldstone phase is qualitatively different: although some logical information remains globally present, puncturing destroys all correlations with the reference. Concretely, in the Goldstone regime we find

$$I(R{:}Q) \to \frac{1}{3} \qquad \text{but} \qquad I(R{:}Q_{p^c}) \to 0, \tag{4.26}$$

so discarding a finite fraction of the bulk eliminates the accessible information about the logical input. By the contrapositive of Corollary 4.2, this rules out any quasi-local recovery channel in the Goldstone phase.

This collapse is again naturally explained by the loop picture. In the Goldstone phase the boundary mode delocalizes: strands emanating from the corner Majoranas wander throughout the bulk on all length scales, so the logical sector (which boundary pairing is realized) is encoded in extended bulk correlations rather than boundary-local data. A simple way to estimate the effect of puncturing is to use the Goldstone scaling of a typical corner strand. In this phase, the typical strand length grows superlinearly with system size, $\ell_{\text{strand}} \sim L^{1+x}$ for some $x > 0$ (with $x = 1$ up to logarithmic corrections, consistent with the $d_f = 2$ fractal scaling reported in Ref. [160]). Now partition the strand into $O(L)$-long coarse-grained segments. Because the puncture is a fixed macroscopic interior window, each

such segment has a finite avoidance probability $p_a < 1$ (equivalently, a finite $O(1)$ chance to intersect Q_p). Since there are $\sim \ell_{\text{strand}}/L \sim L^x$ such segments, the probability that the strand avoids the puncture entirely is at most $\Pr[\text{avoid } Q_p] \sim p_a^{L^x} \to 0$. Thus a bulk puncture of linear size $\Theta(L)$ is intersected with probability approaching one by the strands that encode the boundary pairing, and tracing out Q_p erases that pairing information with probability approaching one, forcing $I(R{:}Q_{p^c}) \to 0$.

Finally, the transition extracted from $I(R{:}Q_{p^c})$ tracks the loop-model transition, showing that the loss of quasi-local recoverability is controlled by the same phase structure that governs the loop statistics.

4.4.4 Decoding transition: global vs local

In this subsection we formulate the short-loop–Goldstone transition as a transition in the *circuit complexity* of recovery: while the logical information can remain globally recoverable in an extended parameter regime, the ability to decode using a quasi-local circuit can fail.

For quasi-local recovery there is a known *sufficient* condition based on the Markov length: if the noise can be embedded into an evolution generated by a local Lindbladian and the Markov length remains finite along the evolution (including for partial-layer intermediate states), then the evolution is reversible by a quasi-local recovery map [193]. In addition, we have introduced a *necessary* condition based on the puncturing criterion: a persistent drop of $I(R{:}Q_{p^c})$ under subdistance punctures rules out any quasi-local recovery by the contrapositive of Theorem 4.1. Taken together, these criteria strongly constrain when quasi-local decoding is possible.

However, these conditions do not by themselves complete the argument in our setting. The sufficiency theorem of Ref. [193] requires that the channel (or an appropriate path connecting the noisy and reference states) arises from a local Lindbladian evolution, and it is not *a priori* clear that our setting admits such a local Lindbladian connecting $\rho_{0,0}$ to $\rho_{p,q}$. We do not rule out this possibility; therefore, the finiteness of the Markov length should be regarded as suggestive rather than decisive evidence for quasi-local recoverability. For this reason, we construct an explicit quasi-local recovery map to establish that the short-loop phase admits quasi-local decoding.

Combining the explicit quasi-local decoder in the short-loop phase with (i) the optimal global decoder and (ii) the puncturing-based obstruction to quasi-local decoding in the Goldstone phase from the previous section, we conclude that the short-loop–Goldstone transition can be interpreted as a transition in the circuit complexity of recovery.

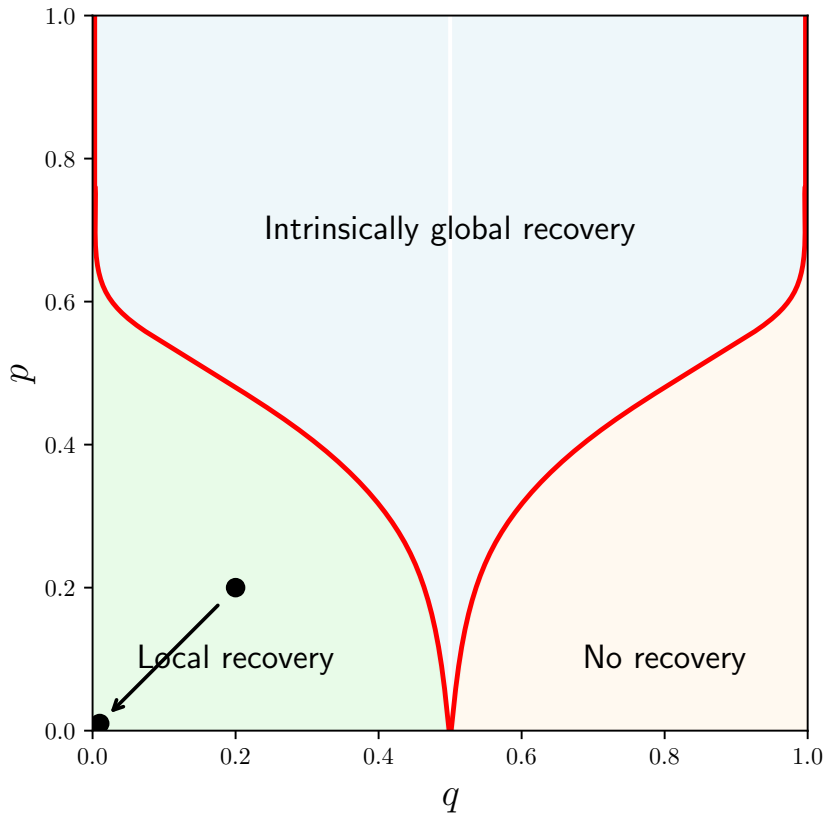


Figure 4.12: The phase diagram from the perspective of recoverability.

Goldstone phase: decoding is intrinsically global

Optimal global decoder. Because the flagged channel is a classical mixture of nonselective single-qubit Pauli measurements, each flag pattern (trajectory) f induces an effective logical nonselective measurement $\bar{\sigma}(f) \in \{\bar{X}, \bar{Y}, \bar{Z}\}$. Only trajectories in the $\bar{X}$ sector can retain any correlation with the reference R; when $\bar{\sigma}(f) \in \{\bar{Y}, \bar{Z}\}$ the output is independent of the encoded $\bar{X}$ bit, so no decoder can do better than random guessing. An optimal decoder therefore has the following form: (i) infer $\bar{\sigma}(f)$ from the flag record, (ii) if $\bar{\sigma}(f) \neq \bar{X}$ output a random guess, and (iii) if $\bar{\sigma}(f) = \bar{X}$ apply a recovery that *canonicalizes* the conditional ensemble σ_f to a fixed reference-point memory state while preserving the logical $\bar{X}$ value.

Concretely, conditioned on $\bar{\sigma}(f) = \bar{X}$, the flags determine at least one viable $\bar{X}$ representative; the recovery uses this information to map σ_f to the clean memory ensemble

(the p=q=0 $\bar{X}$ memory) while carrying the logical bit along, so that the output has the same correlation with R as the clean memory state. A simple choice for such an operation is to first transfer the logical information from an $\bar{X}$ representative to one of the qubits, say qubit k. (This can be done, e.g., by a measure $\bar{X}$ – reset X_k – conditional Z_k – discard measurement outcome channel.) We then reset $X_{j\neq k} \to 1$, transfer the information from X_k to a $q = p = 0$ logical $\bar{X}$ (with k not in its support), reset $X_k \to 1$, and dephase all surface-code Z-stabilizers. This gives the $q = p = 0$ state with the correct logical information.

In the Goldstone phase the logical sector is delocalized, and in the large-system limit the three sectors occur with equal probability, $\Pr\big[\bar{\sigma}(f) = \bar{X}\big] \to 1/3$. Hence even the globally decoder succeeds only with probability $p^{\star}_{\text{succ}} = \frac{1}{3} \cdot 1 + \frac{2}{3} \cdot \frac{1}{2} = \frac{2}{3}$, equivalently $I(R{:}Q) \to 1/3$. The decoder above is optimal in the sense that, whenever correlations with the reference are preserved, it recovers the logical information up to the information-theoretic bound; however, it cannot be implemented by any quasi-local circuit in the Goldstone phase. Operationally, this is exactly what our puncturing diagnostic captures: while the full mutual information remains nonzero, $I(R{:}Q) \to 1/3$, the punctured mutual information vanishes, $I(R{:}Q_{p^c}) \to 0$ for subdistance punctures (Sec. 4.4.3), ruling out quasi-local recovery by the contrapositive of Corollary 4.2.

Short-loop phase: optimal decoder admits a quasi-local implementation

The recovery described above also applies in the short-loop phase: in the $\bar{X}$-pinned regime one has $\Pr\big[\bar{\sigma}(f) = \bar{X}\big] \to 1$, so the same optimal strategy succeeds with probability asymptotically approaching 1. Here, however, our goal is sharper: we aim to establish the existence of a *quasi-local* decoder in the short-loop phase. To do so, we first present a global recovery that is tailored to the short-loop regime and then show how it can be truncated to a quasi-local circuit.

Global decoder via loop/stabilizer updates. Fix a flag configuration f with $\bar{\sigma}(f) = \bar{X}$ and write the corresponding conditional ensemble as σ_f, so that $\rho_{p,q} = \sum_f p(f)\,|f\rangle\langle f| \otimes \sigma_f$ as in Eq. (4.7). The recovery step can be organized as a site by site procedure that turns (Y/Z) sites one by one in to (X) sites. We express this in terms of single-site updates of the form

$$\sigma_f \longmapsto \sigma_{f^{(i\to X)}}, \tag{4.27}$$

where $f^{(i\to X)}$ denotes the flag pattern obtained from f by replacing the local rewiring at site i by the X-rewiring, leaving all other flags unchanged. In the loop picture, such an

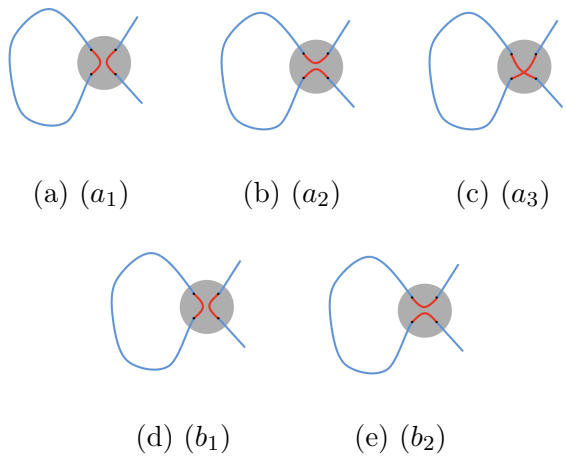
update changes only the reconnection of the four half-edges incident on i, so it affects only the one or two loop constraints carried by the strands passing through i; all loop-parity constraints supported on loops that avoid i are identical before and after the update. Iterating (4.27) over all sites yields a global recovery map that drives σ_f to the canonical $p = q = 0$ ensemble while preserving the logical $\bar{X}$ value. To show this, it suffices to verify that our CPTP map updates the state *consistently* with the updated loop data: after each local rewiring, the resulting state is stabilized by the stabilizers inferred from the correspondingly updated loops. Since the full stabilizer group for each configuration is generated by its loops, enforcing this consistency at every step guarantees that the overall iteration implements the desired transformation from σ_f to $\sigma_{f^{(i\to X)}}$ while maintaining the logical sector.

Here we explain how such a *consistent* loop and stabilizer update works. Fix a site i and consider the four half-edges incident on i. Each half-edge is either paired to another half-edge at i as part of a loop connection, or it is connected (through a strand) to a dangling boundary Majorana, i.e., it participates in a logical strand. For the update rule, the only required input is which *pairing* of half-edges is realized: once we identify one connected pair, the remaining two half-edges may be in either of the two possibilities above, and the update proceeds the same way.

Suppose, for definiteness, that two of the half-edges are connected by a loop segment, while the other two are in an a priori unknown configuration (as Fig. 4.13). The local update then replaces the non-canonical (Y/Z) pairing at i by the canonical X-pairing. Operationally, this means (i) changing the local rewiring at i and (ii) performing the corresponding stabilizer update so that the set of loop-parity constraints remains consistent. Given i, one first identifies the relevant strand connectivity of the incident half-edges, as illustrated in Fig. 4.13(a_1)–(a_3). Depending on whether the site is of Y-type or Z-type and on the sublattice, there are three possible loop patterns. The update replaces these by the appropriate X-type pattern, Fig. 4.13(b_1) or (b_2), again depending on the sublattice. In terms of loop constraints, this replacement either *merges* loops (e.g., transitioning from configuration (a_1) to (b_2), where a loop on one side is reconnected into two strands on the other side), or *splits* a loop (e.g., transitioning from (a_3) to (b_1), where a single connection is reconfigured into a loop plus two separate segments).

We now describe how the stabilizer (and $\bar{X}$-representative) constraints are updated under a single X_i-type re-routing. There are two cases.

Splitting update ($P \mapsto (Q, S)$). Assume the loop configuration shows that the flagged qubit i routes two segments of the *same* strand (for instance, a configuration of the type in Fig. 4.13(a_2) or (a_3)). Let P be the operator for that strand (either a stabilizer loop



(a) (a_1) (b) (a_2) (c) (a_3)

(d) (b_1) (e) (b_2)

Figure 4.13: Local loop rewirings used in the single-site update (4.27). Top row: the three possible non-canonical (Y/Z) routings of the four half-edges incident on a flagged site i (for the different choices of Pauli type (Y/Z) and sublattice), labeled (a_1)–(a_3). Bottom row: the corresponding canonical X-type routings, labeled (b_1) and (b_2). Replacing (a_i) by the appropriate (b_j) changes only the local pairing at i and therefore updates only the loop-parity constraints on the one or two strands passing through i: depending on the incident connectivity, the move either splits a strand into two or merges two strands into one, while leaving all constraints supported on loops that avoid i unchanged.

operator or a $\bar{X}$ representative). Since the strand visits i twice, P does not have i in its support. After performing the X_i-type re-routing, the single strand is resolved into two strands (for instance, a configuration of the type in Fig. 4.13(b_1)) with operators Q and S, and both Q and S have site i in their support. These operators satisfy $P = QS$. As P corresponds to a single strand, at least one factor, say S, is a closed loop. The stabilizer update is implemented by measuring this closed-loop operator S. If the outcome is -1, apply Z_i so that the post-measurement state is stabilized by S with the right eigenvalue. This introduces S into the stabilizer group. Since $P = QS$ and P was already a valid constraint (either a stabilizer or a $\bar{X}$ representative), fixing S simultaneously fixes Q via $Q = PS$: if P was a stabilizer then Q is a stabilizer, while if P was a $\bar{X}$ representative then Q becomes the updated $\bar{X}$ representative. In this way, the single constraint P is replaced by the pair (Q, S), realizing the splitting move.

Merging update ($(A, B) \mapsto AB$). Assume instead that the loop configuration shows that the flagged qubit i routes segments from *different* strands (for instance, a configuration of the type in Fig. 4.13(a_1)). Let A and B be the operators for these two strands (each of A, B

is either a stabilizer loop operator or a $\bar{X}$ representative). Under the X_i-type re-routing, the two strands reconnect into a single strand with operator AB, and this new strand has no support on site i. The desired update therefore keeps AB as the constraint (a stabilizer, or a $\bar{X}$ representative if exactly one of A, B was one) while removing A and B as separate constraints.

A convenient way to implement this merge at the stabilizer level is by a complete X_i dephasing channel, i.e., a nonselective measurement of X_i (measure X_i and discard the outcome). Because A and B each have the site i in their support, they anticommute with the X_i measurement and are removed by the dephasing. In contrast, the product AB has no support on i, so it commutes with the dephasing and remains as a valid constraint. Operationally, this replaces the two constraints A and B by their product AB, which is exactly the merging update.

The process explained above assumes either A or B stabilizers are loop like, however there is the exception when *both* A and B are $\bar{X}$ representatives. In the short-loop phase this event is exponentially suppressed because viable $\bar{X}$ representatives are localized near the boundaries; nevertheless, the global decoder can handle it by first moving the logical information onto site i (or another nearby qubit), recording i for later logical retrieval, and then proceeding with the same local merge rule.

Quasi-local truncation in the short-loop phase. The intuition behind the truncation is that performing the site-by-site loop deformation requires measuring an appropriate *closed-loop* operator (one of S, A, or B in the merge/split description above) and applying a local Pauli correction if needed. In the short-loop phase, loops are localized: with high probability the relevant closed-loop segment needed to identify and complete the update is contained in a bounded neighborhood. This suggests that it suffices to access only a local patch to decide the update rule and perform the corresponding measurement/correction.

To make this intuition precise, we explicitly introduce a stability assumption for the short-loop phase under local loop rewiring. Throughout the qubit-by-qubit decoding procedure, the system realizes an inhomogeneous loop ensemble in which some regions have already been converted to the canonical $(p, q) = (0, 0)$ (X-connected) connectivity, while the remaining regions are still drawn from parameters (p, q) in the short-loop phase. We assume that such non uniform instances do not induce delocalization: the loop statistics remain exponentially localized. Equivalently, converting a subregion to $(p, q) = (0, 0)$ does not generate a growing length scale beyond ξ_{loop}. This assumption is justified since (p, q) and $(0, 0)$ lie in the same short-loop phase: a spatially non-uniform instance obtained by converting a subset of regions to $(0, 0)$ connectivity is therefore still expected to remain

within the same phase, and hence to exhibit short-loop (exponentially localized) statistics controlled by ξ_{loop}.

With this assumption in hand, we now show that the global decoder can be truncated to a quasi-local circuit. Concretely we consider square tiles of linear size ℓ. We divide each tile R into a central square C_R of linear size $(1-a)\ell$ and its complement $\bar{C}_R$, an annulus of thickness $a\ell$, to serve as a buffer region. Recovery proceeds as follows: for each R: (i) we read all flags, (ii) perform the above code deformation in C_R if R has the necessary information and abort otherwise. Provided no R is aborted, covering the system requires a constant number of layers: since the C_R from tiling the system cover area $(1-a)^2L^2$, we need $1/(1-a)^2$ layers of tiles to cover the whole system.

The decoder aborts if, for a nonzero flag in C_R, no closed-loop segments (i.e., S, A or B) of the loop-deformation steps lie fully in R. The probability of this event is upper bounded by the probability p_R that there exists at least one $\bar{C}_R$-spanning loopstrand. By Markov's inequality, this satisfies $p_R \leq \langle N_R \rangle_f$, where N_R is the number of loops spanning $\bar{C}_R$. Here we assume this average quantity has the same statistics everywhere; In particular, loop statistics may differ near the boundary by at most an $O(1)$ factor in the effective loop correlation length ξ_{loop}, and for brevity we ignore this boundary–bulk difference and treat ξ_{loop} as uniform throughout. Consequently, any bulk bound on $\langle N_R \rangle_f$ applies equally well near the boundary. In the short-loop phase, $\langle N_R \rangle_f$ in the bulk is upper bounded by the spanning number (the number of strands spanning $\bar{C}_R$), since each $\bar{C}_R$-spanning loop contributes at least one $\bar{C}_R$-spanning strand. Using the known behavior of the spanning number in the bulk, we obtain exponential decay in the bulk, and by the same homogeneity assumption we apply the same bound near the boundary. For $a\ell \gg \xi_{\text{loop}}$, we have $\langle N_R \rangle_f \propto \text{poly}(\ell) \exp(-a\ell/\xi_{\text{loop}})$. We can hence estimate [4] how $a\ell$ must scale to suppress the failure rate P_{fail} with L, i.e., the probability that the decoder aborts for at least one region R. By a union bound and up to constant factors,

$$P_{\text{fail}} \lesssim L^2 p_R \lesssim L^2 \,\text{poly}(\ell) \exp(-a\ell/\xi_{\text{loop}}). \tag{4.28}$$

Hence, requiring $a\ell/\xi_{\text{loop}} > \log L^x$ with $x > 2$ is sufficient for P_{fail} to be suppressed as $L \to \infty$. In particular, choosing $\ell = (L)$ gives a *polylogarithmic range* decoder (the recovery map acts only within tiles of diameter $O(\ell)$), and the circuit depth is also (L) because within each tile the loop deformation updates on the sites in C_R must be performed sequentially

[4] In addition to the abort mechanism captured by p_R, one can consider a distinct logical failure mode in which the lower and upper logical $\bar{X}$ strands both reside within the same tile. In the short-loop phase this contribution is still suppressed as $\exp(-L/\xi_{\text{loop}})$, and is therefore negligible compared to the dominant contribution estimated in Eq 4.28.)

as the required loop measurements and stabilizer updates depend on the outcomes and on previously updated local pairings.

Finally, we emphasize that the above quasi-local truncation argument is intended only for the $\bar{X}$-pinned short-loop phase. It does not apply in the other phases, nor does it apply to the rare trajectories (even within the short-loop regime) where the effective logical measurement is $\bar{Z}$ or $\bar{Y}$: in those cases global recovery already fails, so one should not expect the quasi-local decoder to succeed. Equivalently, such trajectories are not in the $\bar{X}$-pinned short-loop sector—they either already contain long strands/loops, or effectively lie in the $\bar{Z}$-pinned short-loop phase. In either case, inserting the $(p, q) = (0, 0)$ connectivity into a patch creates an interface between distinct phases, which generically produces long strands. Consequently, the stability assumption used in the truncation argument breaks down.

4.5 Discussion

We have analyzed an ensemble obtained from the surface code under heralded random Pauli noise and shown that it realizes an extended mixed-state critical phase with short-range local correlations but divergent Markov length, in particular polylogarithmic decay of CMI. In terms of error correction properties, we found using coherent information that the critical phase retains a partial logical memory that is only nonlocally recoverable. To rule out quasi-local decoding, we introduced a punctured version of coherent information. In contrast, in the phase corresponding to short loops in which logical memory is preserved, we proposed an explicit quasi-local decoder.

It would be interesting to generalize the construction to other low-density parity check (LDPC) codes and lattice geometries, where the underlying loop or network models can exhibit different symmetry classes and possibly different types of critical phases. One could explore dynamical versions of this setup, in which coherent noise and measurements act repeatedly, to study the interplay between mixed-state criticality and measurement-induced transitions. Furthermore, it would be useful to understand whether the critical phase in this work is robust to certain perturbations, and relatedly, if it can be understood in terms of strong-to-weak spontaneous symmetry breaking (SWSSB) [130, 133, 189].

While the current work focused on critical mixed state phases that are separable, it would be interesting to find an example of a *quantum* critical mixed-state phase, characterized not only by sub-exponential decay of CMI but also long-range entanglement. Verifying this possibility [97], and its implications for quantum memory, is a work in progress.

Chapter 5

Symmetry enforces entanglement at high temperatures

This chapter investigates whether symmetry constraints can prevent the sudden death of entanglement in thermal states of locally interacting many-body systems. We prove that strongly symmetric thermal states in the canonical ensemble, for generic Hamiltonians with on-site Abelian symmetries, remain entangled with nonzero entanglement negativity at arbitrarily high temperatures, under mild conditions on the symmetry action and the charge sector. Our results extend to weakly symmetric thermal states in the Gibbs ensemble under superselection rules, which restrict state decompositions to be symmetric. In particular, we show that fermionic Gibbs states evade sudden death of entanglement and retain nonzero fermionic negativity at high temperatures, proving along the way several existing conjectures about fermionic entanglement. These findings demonstrate that global symmetry correlations can preserve quantum entanglement despite thermal decoherence, and provide new insight into the interplay between symmetry and quantum information in thermal equilibrium.

5.1 Introduction

Many-body systems equilibrate to their Gibbs distribution at finite temperatures by exchanging energy and information with a bath. The resulting state $\rho_\beta \propto e^{-\beta H}$ for system Hamiltonian H at inverse temperature $\beta = T^{-1}$ is a mixed state that maximizes entropy under energy constraints, reflecting an observer's ignorance of information leaked to the

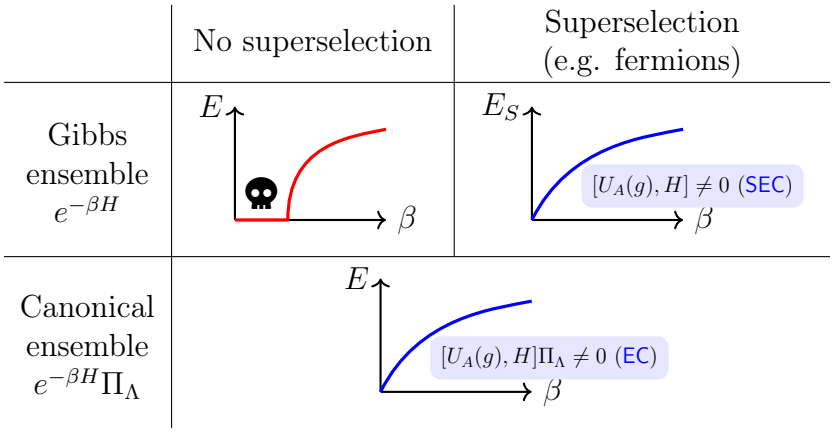


Table 5.1: Typical behavior of entanglement measures (E, E_S) as functions of the inverse temperature $\beta = T^{-1}$ for the Gibbs and canonical ensembles in the cases with and without superselection. E (E_S) is a measure of (symmetric) entanglement, such as (fermionic) negativity. There is sudden death of entanglement for the Gibbs ensemble with no superselection rule if H is local [12], while there is not for the other cases if their respective entangling conditions (EC and SEC) are satisfied. For an example of the above, see Fig. 5.2.

bath. For physical, local quantum Hamiltonians, the Gibbs state has structural properties at finite non-zero temperatures that reflect its low complexity, such as area law of mutual information [235] and of entanglement [120], and local Markov property [43, 119, 104]. Importantly, it was recently proven [12] that Gibbs states of local Hamiltonians undergo "sudden death of entanglement" (SDOE) at sufficiently high but finite temperatures in the thermodynamic limit, becoming exactly representable by a classical distribution of pure product states. This result confirmed evidence from previously studied few-body and solvable models [10, 70, 5, 237, 9, 208, 84, 170, 82] (See top left plot of Table 5.1).

Meanwhile, recent developments in mixed-state phases of matter have confirmed the essential role of symmetries [130, 189, 133, 240, 151, 152, 131, 236, 150] and uncovered novel connections to their global entanglement properties [156, 139, 187, 132, 134, 144, 231, 137, 229, 99]. In this Letter, we prove that symmetry restrictions on thermal states generically prevent SDOE. For that, we consider a symmetry group G, and a Hamiltonian H that is symmetric under G: $\forall g \in G, [U(g), H] = 0$, where $U : G \to U(\mathcal{H})$ is a symmetry representation of G in the total Hilbert space $\mathcal{H}$. Gibbs states of local Hamiltonians H that are symmetric also undergo SDOE, so we must impose a stronger symmetry constraint on the Gibbs state to see meaningful effects on entanglement. We consider two physically

motivated ways of doing that:

One way is to impose **strong symmetry** in the thermal ensemble $\rho_\beta = \sum_i p_i |\psi_i\rangle\langle\psi_i|$ by limiting the mixture only to states $|\psi\rangle$ that have a fixed symmetry charge $\Lambda : G \to U(1)$, i.e. $U(g)|\psi\rangle = \Lambda(g)|\psi\rangle$. This constraint results in the *canonical ensemble* $\rho_{\beta,\Lambda} \propto e^{-\beta H}\Pi_\Lambda$, where Π_Λ projects onto the charge sector Λ. By definition, $\rho_{\beta,\Lambda}$ is *strongly symmetric* under G, satisfying $U(g)\rho_{\beta,\Lambda} = \Lambda(g)\rho_{\beta,\Lambda}$, which is a stronger requirement than the usual *weak symmetry* condition $U(g)\rho U(g)^\dagger = \rho$ [38, 7, 6]. The canonical ensemble naturally arises when an initial strongly symmetric system is put into contact with a thermal bath in a way that preserves the symmetry of the system alone (See Appendix D.1 for a proof).

Alternatively, we can impose a **superselection rule** by considering a restricted quantum theory consisting only of states that are (weakly) symmetric under G [15, 74, 40]. This restriction can either be fundamental to physical reality, as in the case of the fermion parity superselection rule [91], or be put by hand, as in real quantum theory [216, 39, 238]. Crucially, the superselection rule alters the definition of entanglement itself [221, 199, 201, 16, 166]. Namely, we call a bipartite state ρ_{AB} *symmetrically separable* if it is the mixture of tensor product states $\rho_A \otimes \rho_B$ whose parts, ρ_A and ρ_B, are also valid states within the superselected theory, and thus weakly symmetric by themselves. This is generally stronger than the usual separability condition [234], since there can be separable states that are not symmetrically separable.

These two symmetry conditions on ρ_β have similar entanglement properties in the following sense: ρ_β is symmetrically entangled (i.e. *not* symmetrically separable) if, and only if, the canonical ensemble $\rho_{\beta,\Lambda}$ is also entangled for some charge Λ (See Appendix D.3 for a simple proof). Then, we are able to prove the persistence of entanglement[1] for both cases at high temperatures (See Fig. 5.1):

Theorem 5.1 (Persistence of entanglement: informal)**.** *For Abelian on-site symmetries and generic symmetric local Hamiltonians at sufficiently high temperatures,*

1. *the Gibbs ensemble ρ_β is symmetrically entangled, and*
2. *the canonical ensemble $\rho_{\beta,\Lambda}$ is entangled with non-zero entanglement negativity, for almost all symmetry actions and irreps Λ in the thermodynamic limit.*

The second part of the theorem above is stronger than the first in two ways: first, we provide necessary and sufficient conditions for $\rho_{\beta,\Lambda}$ to be entangled, with Λ being arbitrarily

[1] Not to be confused with the robustness measure of multipartite entanglement called "persistency of entanglement" [34, 37].

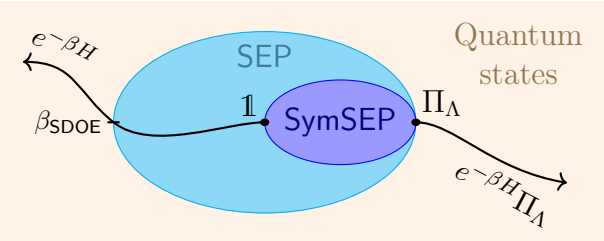


Figure 5.1: Schematic paths of generic many-body thermal states in relation to the set of (symmetrically) separable states SEP (SymSEP). The Gibbs ensemble $\rho_\beta \propto e^{-\beta H}$ is separable from $\beta = 0$ to β_{SDOE} due to the sudden death of entanglement, but it becomes symmetrically entangled as soon as $\beta > 0$. Similarly, the canonical ensemble $\rho_{\beta,\Lambda} \propto e^{-\beta H}\Pi_\Lambda$ is also entangled for $\beta > 0$. For a more accurate representation of SEP and SymSEP in a slice of the space of quantum states, see Fig. D.1 in Appendix D.2.

chosen; and, second, $\rho_{\beta,\Lambda}$ is not only entangled, but has nonzero negativity, which is a computable measure of entanglement [171, 94, 223]. The phrase "for almost all symmetry actions and irreps" excludes the condition of *semiuniformity* (Def. 5.3), which is absent in most examples of interest, such as finite G at sufficiently large system size.

We only consider *Abelian* symmetry groups with *on-site* action because their Gibbs states are fully symmetrically separable at infinite temperature. In contrast, it has been found that non-Abelian symmetries [156, 139, 187] and anomalous non-on-site symmetries [132, 144, 231, 137, 229, 99] can guarantee entanglement even at infinite temperature, so it would not be as surprising if the entanglement of their symmetric thermal ensembles persists at high temperatures.

Recent works report similar findings on the persistence of entanglement [73] and negativity [109] in canonical ensembles at high temperatures. Our work provides significant generalization through distinct proof techniques and complete characterization of symmetries and irreps with persistent entanglement. Furthermore, we unify the "no SDOE" result for canonical ensembles with the persistence of symmetric entanglement in thermal states under superselection rule, which were previously noted for fermionic systems [16, 18, 206, 47, 179]. More specifically for fermions, we are also able to show an equivalent version of Theorem 5.1 where the entanglement negativity measure used is the appropriate one for fermionic systems [207, 205].

5.2 Setup and notation

We consider a multipartite Hilbert space with $N < \infty$ identical sites, $\mathcal{H} = \bigotimes_{i=1}^{N} \mathcal{H}_i$, $\forall i,j. \mathcal{H}_i \simeq \mathcal{H}_j$. The symmetry group G is Abelian and acts as a on-site unitary $U : G \to \mathcal{L}(\mathcal{H})$, $U = \bigotimes_i u_i$ in a non-trivial way on each site, i.e. $\exists g \in G, u(g) \not\propto \mathbb{1}$. Sometimes, we will specialize to homogeneous symmetry actions $U = u^{\otimes N}$.

From the symmetry action of G on the Hilbert space $\mathcal{H}$, it is decomposed into $\mathcal{H} = \bigoplus_\lambda V_\lambda^{\otimes m_\lambda}$, where the direct sum is over all irreps λ of G, with m_λ being the multiplicities. Since G is Abelian, each irrep space V_λ is one-dimensional, and so the irreps are identified with their characters, $\lambda : G \to U(1)$. By Schur's lemma, a weakly symmetric state ρ can be decomposed into strongly symmetric block diagonal components, $\rho = \sum_\lambda q_\lambda \rho_\lambda$, where $\rho_\lambda = \rho\Pi_\lambda/q_\lambda$ with $\Pi_\lambda = \sum_{g\in G} \overline{\lambda(g)} U(g)$ being the projector onto the *isotypic component* $V_\lambda^{\otimes m_\lambda}$, and $q_\lambda = \mathrm{Tr}[\rho\Pi_\lambda]$. We shall call $\sum_\lambda q_\lambda \rho_\lambda$ the *isotypic decomposition* of ρ.

A state ρ is separable ($\rho \in \mathsf{SEP}$) with respect to a bipartition $A|B$ if it can be represented as a convex mixture of tensor product states, $\rho = \sum_i p_i \rho_{A,i} \otimes \rho_{B,i}$ [234]. If a separable state ρ is weakly symmetric under a bipartite symmetry representation $U = U_A \otimes U_B$, and each $\rho_{A,i}$ and $\rho_{B,i}$ is weakly symmetric under U_A and U_B respectively, then we say ρ is *symmetrically separable* ($\rho \in \mathsf{SymSEP}$) [2].

5.3 Persistence of entanglement

Now consider the Gibbs state $\rho_\beta \propto e^{-\beta H}$. If H is local, ρ_β becomes fully separable at high temperatures. We will now give conditions under which not only symmetry constraints *prevent* SDOE, but also lead to the *persistence* of entanglement (POE) at high temperature, by which we mean that there exists a $\beta^* > 0$ such that ρ_β is symmetrically entangled if $\beta < \beta^*$, with an analogous statement for $\rho_{\beta,\Lambda}$.

We proceed by assuming that the entanglement of ρ_β does *not* persist. That is equivalent to ρ_β being symmetrically separable for all β in a decreasing sequence $(\beta_n)_{n\in\mathbb{N}}$ converging to zero: $\rho_\beta = \sum_i p_i \rho_{\beta,A,i} \otimes \rho_{\beta,B,i}$ for a bipartition $A|B$. Since the $\rho_{\beta,A,i}$ states are individually symmetric, we have $\forall g \in G, [U_A(g), \rho_\beta] = \sum_i p_i [U_A(g), \rho_{\beta,A,i}] \otimes \rho_{\beta,B,i} = 0$. Since each matrix element of $[U_A(g), \rho_\beta]$ is a holomorphic function of β that is zero on

[2] For on-site symmetries this is in fact equivalent to the following definition of symmetric separability [16]: ρ is symmetrically separable for a symmetry U if there exists a decomposition $\rho = \sum_i p_i \rho_i$, such that $[\rho_i, U] = 0$ and ρ_i is a tensor product state.

the set $\{\beta_n\}$ with the accumulation point $\beta = 0$, it is identically zero for all $\beta \in \mathbb{C}$. In particular, its derivative at $\beta = 0$ must also be zero, which leads to a necessary condition for symmetric separability at high temperatures: $\forall g \in G,\ [U_A(g), H] = 0$.

By the converse of the reasoning above, we arrive at a sufficient condition for the Gibbs state to be symmetrically entangled at arbitrarily high temperatures $\beta \in (0, \beta^*)$,

$$\forall g \in G,\ [U_A(g), H] \neq 0, \qquad \text{[Symmetric entangling condition (SEC)]} \tag{5.1}$$

By the correspondence between the symmetric separability of ρ_β and the separability of $\rho_{\beta,\Lambda}$, SEC also implies the existence of at least one irrep Λ whose canonical state exhibits POE. Beyond this, following the same argument above for $\rho_{\beta,\Lambda} \propto e^{-\beta H}\Pi_\Lambda$, we can reach a different sufficient condition for POE for *any given* Λ,

$$\forall g \in G,\ [U_A(g), H]\Pi_\Lambda \neq 0. \qquad \text{[Entangling condition (EC)]} \tag{5.2}$$

Combining these results, we arrive at the following theorem, which is summarized in Table 5.1:

Theorem 5.2 (Persistence of entanglement)**.** *Given an on-site Abelian symmetry,*

1. *For a symmetric Hamiltonian H that satisfies SEC, the Gibbs ensemble is symmetrically entangled, and $\exists$ an irrep Λ such that the canonical ensemble is entangled at arbitrarily high temperatures.*

2. *For a symmetric Hamiltonian H and an irrep Λ that satisfies EC, the canonical ensemble is entangled at arbitrarily high temperatures.*

5.4 Genericness and local indistinguishability

We now intend to prove that SEC and EC are generically satisfied. The crucial property of these conditions is that they are linear in H; for any H that violates SEC/EC, adding a perturbation V that satisfies SEC/EC ensures that $H + \epsilon V$ also satisfies the condition for all $\epsilon \neq 0$. Thus, proving the existence of entangling perturbations suffices to prove that fine-tuning H is required to violate SEC/EC.

For example, consider a qubit chain with $G = \mathbb{Z}_2$ symmetry generated by $\bigotimes_i Z_i$. A fine-tuned Hamiltonian H_Z consisting solely of Z terms yields separable thermal states ρ_β

and $\rho_{\beta,\lambda}$. However, the operators X_i and Y_i carry charge -1, so $V_{ij} = X_iX_j$ or Y_iY_j is an entangling perturbation when $i \in A$ and $j \in B$, making $H_Z + \epsilon V_{ij}$ satisfy both EC and SEC, even for infinitesimal ϵ.

For other symmetry groups, an entangling V that satisfies SEC can be similarly constructed by pairing a charged operator in A with an oppositely charged one in B, so that $[U(g), V] = 0$ but $[U_A(g), V] \neq 0$. The EC, however, is harder to establish in general, due to the projector Π_Λ. In fact, there are some charge sectors for which EC is not satisfiable: e.g. the maximal or minimal charge irrep of a $U(1)$ action on a system of qubits as $\theta \mapsto e^{i\theta \sum_i Z_i}$ leads to an unentangled canonical ensemble, as the projector is rank one, negating EC regardless of V.

We can avoid such scenarios if we specialize to *finite* Abelian symmetries with homogeneous symmetry action, $U = u^{\otimes N}$, which allows us to use local indistinguishability of symmetric subspaces to justify the existence of V satisfying both EC and SEC. Consider a large tripartition $A|B|C$, where B is a buffer region that separates A and C, taken to be large enough such that there are no local Hamiltonian terms in H supported on AC. Since Π_Λ is maximally mixed with just a *global* charge constraint, we expect $\mathrm{Tr}_C \Pi_\Lambda \propto \mathbb{1}_{AB}$, which we can quantitatively establish for finite Abelian groups (See Appendix D.4). This implies that tracing out a sufficiently large region disjoint from the support of V and A reduces $[U_A(g), V]\Pi_\Lambda = 0$ to $[U_A(g), V] \approx 0$, with equality in the thermodynamic limit. In fact, for finite Abelian groups, we can further show that SEC and EC are actually equivalent for large enough regions in the thermodynamic limit (See Appendix D.5).

Not only is $\rho_{\beta,\Lambda}$ entangled when EC is satisfied, but it is also locally equivalent to ρ_β. We prove this for k-local Hamiltonians symmetric under finite G at high temperatures, in Appendix D.4.

5.5 Classification of irreps

Using local indistinguishability to construct entangling perturbations has two limitations: it assumes large system sizes, and its rigorous formulation requires finite Abelian groups. As already seen, infinite groups like $U(1)$ have irreps that are not only locally distinguishable to identity, but also lead to unentangled canonical ensembles. We now present a construction of entangling perturbations that goes beyond the local indistinguishability argument, using the property of semiuniformity of irreps.

Definition 5.3. A global irrep Λ of an N-partite system is **uniform** if its projector Π_Λ is a uniform tensor product of an on-site irrep λ: $\Pi_\Lambda = \bigotimes_{i=1}^N \Pi_\lambda^{(i)}$. Λ is **semiuniform** if its

projector Π_Λ is a direct sum of uniform tensor products: $\Pi_\Lambda = \sum_\alpha \bigotimes_{i=1}^N \Pi_{\lambda_\alpha}^{(i)}$.

A small-sized example of a semiuniform irrep is the trivial irrep $\Lambda = 1$ of a $\mathbb{Z}_3$ action on $N = 3$ qubits generated by $\tilde{Z}_1\tilde{Z}_2\tilde{Z}_3$, where $\tilde{Z} = \mathrm{diag}(1, e^{2\pi i/3})$. In this case, $\Pi_{\Lambda=1}$ projects onto the subspace generated by the uniform states $|000\rangle$ and $|111\rangle$. An example of a uniform irrep is the maximal or minimal charge irrep of the aforementioned $U(1)$ symmetry on qubits.

By excluding semiuniform irreps, we are able to construct two-body entangling Hamiltonians:

Theorem 5.4. *Consider a multipartite system with homogeneous symmetry action U. If Λ*

1. *is not semiuniform, then for any pair of sites (i, j), there exists a two-body symmetric Hamiltonian $v_{i,j}$ supported on $\{i, j\}$ that entangles any bipartition that separates site i from site j. Moreover, given a connected interaction graph, the uniform nearest-neighbor perturbation $\sum_{\langle i,j\rangle} v_{i,j}$ entangles all bipartitions.*
2. *is semiuniform, then there is no entangling Hamiltonian with support over $N-1$ sites or fewer, and for all $(N-1)$-local Hamiltonians, $\rho_{\beta,\Lambda}$ is fully separable when ρ_β is fully separable.*

We now sketch the construction of the entangling $v_{i,j}$, leaving a full proof to Appendix D.6. If the irrep is not semiuniform, then for any pair of sites i, j there exist at least two local irreps $\lambda' \neq \lambda$. Then let $\sigma_i^+ := |\lambda'\rangle\langle\lambda|_i$ and $\sigma_j^+ := |\lambda'\rangle\langle\lambda|_j$. Denoting their Hermitian conjugates by $\sigma_{i,j}^- := (\sigma_{i,j}^+)^\dagger$, we choose $v_{i,j} := \sigma_i^+\sigma_j^- + \sigma_i^-\sigma_j^+$. It is easy to check that $v_{i,j}$ is symmetric and satisfies the entangling condition Eq. 5.2 if $i \in A$, since it pairs a charged operator in A with an oppositely charged one in B.

In fact, we can also show that semiuniform irreps are rare: if the symmetry is homogeneous and independent of the number of sites, the fraction of semiuniform global irreps goes to zero in the thermodynamic limit as $O(2^d/N)$, where d is the on-site Hilbert space dimension. Furthermore, if G is finite, there are no semiuniform global irreps for $N > |G|$. We prove this in Appendix D.6.

5.6 Persistence of negativity

We can also ask whether $\rho_{\beta,\Lambda}$ remains entangled while exhibiting zero negativity. This question was answered negatively in Ref. [109] for the classical Ising Hamiltonian $H =$

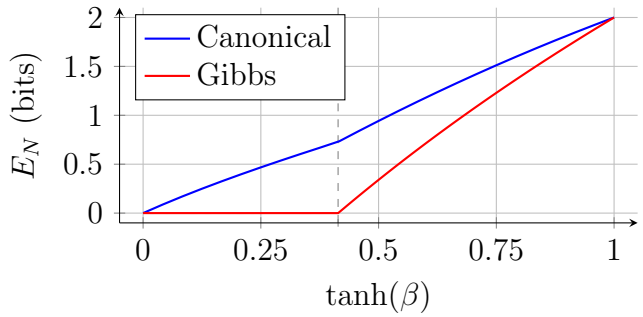


Figure 5.2: Entanglement negativity E_N for the Gibbs and canonical ensembles of the cluster chain Hamiltonian $H = \sum_{i=1}^{N} Z_{i-1}X_iZ_{i+1}$ under $\mathbb{Z}_2$ symmetry generated by $\prod_i X_i$ in the limit $N \to \infty$. Both curves are independent of the entangling interval size $|A| \geq 2$, and the canonical one, of the charge sector.

$\sum_i Z_iZ_{i+1}$ with $\mathbb{Z}_2$ symmetry generated by $\prod_i X_i$. In Appendix D.8, we report the same conclusion for the cluster chain Hamiltonian $H = \sum_i Z_{i-1}X_iZ_{i+1}$ with the same symmetry (See Fig. 5.2).

We generalize these examples by showing that the entanglement condition of Eq. (5.2) also implies nonzero negativity, and thus the absence of bound entanglement with positive partial transpose (PPT)[3]. To establish this, we begin by defining the basis with respect to which the partial transpose is taken.

For the case of on-site Abelian symmetries $U = \bigotimes_i u_i$, we choose to do the partial transpose in a region A with respect to a basis that diagonalizes U_A, such that $\forall g \in G, U_A(g)^{T_A} = U_A(g)$, and thus we would have $\Pi_\Lambda^{T_A} = \Pi_\Lambda$. Now consider the operator $\tilde{\Pi}_\Lambda \, \rho_{\beta,\Lambda}^{T_A} \, \tilde{\Pi}_\Lambda$, with $\tilde{\Pi}_\Lambda := \mathbb{1} - \Pi_\Lambda$. It is traceless, as $\mathrm{Tr}\left(\tilde{\Pi}_\Lambda \, \rho_{\beta,\Lambda}^{T_A} \, \tilde{\Pi}_\Lambda\right) \propto \mathrm{Tr}\left(e^{-\beta H}\Pi_\Lambda \tilde{\Pi}_\Lambda^{T_A}\right) = 0$, and if $\rho_{\beta,\Lambda}$ has PPT, then $\tilde{\Pi}_\Lambda \, \rho_{\beta,\Lambda}^{T_A} \, \tilde{\Pi}_\Lambda$ is also positive semidefinite. Hence, it must be identically zero. Just like before, assuming $\rho_{\beta,\Lambda}$ has PPT in a sequence of β converging to zero implies the derivative of $\tilde{\Pi}_\Lambda \, \rho_{\beta,\Lambda}^{T_A} \, \tilde{\Pi}_\Lambda$ at $\beta = 0$ must vanish, which gives a necessary condition for zero negativity at high temperatures: $(\mathbb{1} - \Pi_\Lambda)(H\Pi_\Lambda)^{T_A}(\mathbb{1} - \Pi_\Lambda) = 0$.

Negating this necessary condition gives a sufficient one for POE with non-zero negativity for a bipartition $A|B$:

$$(\mathbb{1} - \Pi_\Lambda)\,(H\Pi_\Lambda)^{T_A}\,(\mathbb{1} - \Pi_\Lambda) \neq 0. \qquad \text{[Negativity condition (NC)]} \tag{5.3}$$

[3]The logarithmic negativity upper bounds distillable entanglement; however non-zero negativity does not necessarily imply no bound or undistillable entanglement, as there may exist bound entanglement which is nevertheless non-PPT.

In Appendix D.7, we show the equivalence of the entangling and negativity conditions, EC $\iff$ NC. Combining the results so far, we have the general implication for POE in canonical ensembles, and all the results quoted in Theorem 5.1.

5.7 Fermions

We now discuss fermionic systems, where the physically relevant states are always (weakly) symmetric under fermion parity. Consider a local Hilbert space $\mathcal{H}$ and an operator algebra spanned by mutually anticommuting Majorana operators $\{c_j\}$ with $\{c_j, c_k\} = 2\delta_{jk}$, with the fermion parity operator $P \equiv \prod_i c_i$. Physical density operators preserve P and their set is referred to as $\mathcal{S}(\mathcal{H})$. Any (symmetrically) separable state under a bipartition $A|B$ must also preserve the local parity operators P_A, P_B, where $P_{A(B)} \equiv \prod_{i\in A(B)} c_i$ [16]. Our previous results thus imply that the fermion parity superselection rule generically protects entanglement from sudden death at high temperatures. Here, we establish a stronger result: even a computable entanglement measure, the fermionic negativity [207, 205], generically persists at arbitrarily high temperatures for both the Gibbs and the canonical ensembles.

Any density matrix $\rho \in \mathcal{S}(\mathcal{H}_A \otimes \mathcal{H}_B)$ can be expanded in Majorana monomials $a_{p_1} \cdots a_{p_{k_1}} b_{q_1} \cdots b_{q_{k_2}}$, where a_{p_j} and b_{q_j} act on A and B respectively, and $k_1 + k_2$ is even by fermion parity. The fermionic partial transpose $\mathcal{T}_A$ multiplies such a monomial by the phase i^{k_1}. The fermionic negativity, defined as $E_N^f = \log ||\rho^{\mathcal{T}_A}||_1$, is a more sensitive measure of fermionic entanglement [205] than its "bosonic" counterpart [66, 65], which is defined without the phase factor i^{k_1}.

We proceed by contradiction to show that E_N^f persists generically for the *Gibbs ensemble* $\rho_\beta \propto e^{-\beta H}$. Assume ρ_β has zero fermionic negativity, i.e. $||\rho_\beta^{\mathcal{T}_A}||_1 = 1$, for a sequence of β converging to 0. This, combined with the fact that $\mathrm{Tr}\{\rho_\beta^{\mathcal{T}_A}\} = \mathrm{Tr}\{\rho_\beta\} = 1$, implies that $\rho_\beta^{\mathcal{T}_A}$ is Hermitian[4]. By taking the derivative of $\rho_\beta^{\mathcal{T}_A} - (\rho_\beta^{\mathcal{T}_A})^\dagger = 0$ at $\beta = 0$, we have a necessary condition for zero E_N^f at high temperature:

$$H^{\mathcal{T}_A} - (H^{\mathcal{T}_A})^\dagger = 0,$$

whose negation gives a sufficient condition for the persistence of E_N^f,

$$(H^{\mathcal{T}_A})^\dagger \neq H^{\mathcal{T}_A}. \tag{5.4}$$

[4]This is proved in Lemma 2.2 of [143]: $||A||_1 = \mathrm{Tr}\, A$ if and only if $A \geq 0$.

This is satisfied by the two-Majorana term $V = ic_ic_j$, with $i \in A$ and $j \in B$, since $(ic_ic_j)^{\mathcal{T}_A} = -c_ic_j$ is not Hermitian. Note that this entangling condition never occurs for the "bosonic" negativity, since H^{T_A} is always Hermitian. The same entangling perturbation also leads to the persistence of E_N^f for the canonical ensemble with fermion parity, which we explain in Appendix D.9.

5.8 Discussion

Our results imply the persistence of entanglement at high temperatures in two related scenarios: when the thermal state is enforced to be strongly symmetric (canonical ensemble), and when the symmetry is enforced at the level of decomposition of the mixed state into pure states (Gibbs ensemble with superselection). We also discuss the generic nature of this entanglement, showing the existence of local Hamiltonian perturbations which lead to entanglement at arbitrarily high temperatures, as well as a complete classification of charge sectors with POE for homogeneous on-site global Abelian symmetries.

These results can be generalized in various ways. Firstly, the results of the sufficient conditions for POE also extend to Abelian higher-form and subsystem symmetries. In fact, strongly symmetric versions of toric codes (with Abelian one-form symmetries) already show a persistence of negativity, as can be inferred from [84]. A more careful study of higher-form and other symmetries, including antiunitary symmetries, is left for a future study.

Some open questions about the quantification of entanglement at high temperatures remain: our results only imply that a faithful measure of entanglement (and negativity) will remain non-zero at arbitrarily high temperatures for any fixed system size N, but it could vanish as $N \to \infty$. Calculations on stabilizer models, however, suggest that the entanglement negativity does not vanish at non-zero β even in the thermodynamic limit $N \to \infty$, which will be useful to establish more generally. This is also relevant to the question of local detectability of the entanglement in $\rho_{\beta,\Lambda}$ for temperatures past SDOE of ρ_β.

It is also interesting to explore whether quantum resources beyond entanglement, such as magic and Gaussianity [179], undergo sudden death for the canonical ensemble. Lastly, it would be useful to investigate the complexity of preparation of the canonical ensemble, in view of the recent interest in similar questions about the Gibbs ensemble [41, 42, 185, 12].

Chapter 6

Extendibility of Fermionic Gaussian States

This chapter investigates (k_1, k_2)-extendibility of fermionic Gaussian states, a property central to quantum correlations and approximations of separability. We show that these states are (k_1, k_2)-extendible if and only if they admit a fermionic Gaussian extension, yielding a complete covariance-matrix characterization and a simple semidefinite program whose size scales linearly with the number of modes. This provides necessary conditions for arbitrary fermionic states and is sufficient within the Gaussian setting. Our main result is a finite de Finetti-type theorem: we derive trace-norm bounds between (k_1, k_2)-extendible fermionic Gaussian states and separable states, improving previous exponential scaling to linear in the number of modes, together with complementary relative-entropy and squashed- entanglement bounds. For two modes, the upper and lower bounds match at order $1/\sqrt{k_1 k_2}$. Extendibility also provides operational support for one of the different notions of separability in fermionic systems. Finally, for fermionic Gaussian channels, we provide a semidefinite-program criterion for anti-degradability and show that entanglement- breaking channels coincide with replacement channels, implying that there are no nontrivial entanglement-breaking fermionic Gaussian channels.

6.1 Introduction

Quantum entanglement underpins both foundational aspects of quantum theory and its technological applications [96, 167]. Beyond serving as a resource for quantum communication, computation, and metrology [62, 89], its structure reveals fundamental constraints

on how quantum correlations can be distributed [174, 8, 64]. One of the most striking of these constraints is *monogamy* [51, 218, 123, 115]: entanglement cannot be freely shared among multiple parties. This idea is formalized through *extendibility* [59, 60, 138]: A bipartite state ρ^{AB} is said to be (k_1, k_2)-*extendible* if there exists a state $\widetilde{\rho}^{A_1 \ldots A_{k_1} B_1 \ldots B_{k_2}}$ such that (i) it is invariant under permutations of the A-systems and, separately, under permutations of the B-systems, and (ii) for every pair (i, j), the two-body marginal satisfies $\widetilde{\rho}^{A_i B_j} = \rho^{AB}$. Two-sided (k_1, k_2)-extendibility provides a symmetric characterization of entanglement shareability, capturing constraints on both subsystems and offering a more refined operational framework for analyzing bipartite correlations.

A fundamental result is that infinite extendibility is equivalent to separability [69, 61, 219, 48]. This setting naturally relates to finite quantum de Finetti theorems, which characterize permutation-invariant states as approximate convex combinations of independent and identically distributed (i.i.d.) product states [183, 117, 48]. Classical de Finetti results describe infinite extendibility, which imples exact product structure; finite versions, in contract, quantify how close (k_1, k_2)-extendible states are to separable ones, providing operationally meaningful bounds. Although extendibility and de Finetti-type results have been extensively studied for systems with tensor-product Hilbert spaces and commuting observables—such as qubits and bosonic Gaussian states [122, 177]—they remain far less explored in the fermionic setting [117, 118], despite the ubiquity and fundamental importance of fermions across physics.

Fermionic many-body states are foundational in condensed-matter physics and quantum chemistry, underpinning electronic structure and correlated phases [124, 217]. They are likewise important in quantum information, where fermionic modes provide efficient models for simulation, benchmarking, and computation [26, 64, 29, 30]. In the fermionic Gaussian (quasifree) regime these states arise as ground and thermal states of quadratic Hamiltonians, remain closed under quadratic dynamics, and admit efficient correlation-matrix simulation methods [173, 172]; their relevance is illustrated by paradigmatic topological models such as the Kitaev honeycomb lattice and the SSH chain [111, 90]. These systems further motivate work on fermionic resource theories, separability diagnostics, Gaussian channel structure [65, 207, 78, 205].

For fermions, anticommutation relations and superselection rules alter the structure of "local" operations and the notion of entanglement. One can view entanglement either between indistinguishable particles (first quantization) [198, 135] or, more naturally for quantum information, between sets of modes assigned to different spatial regions (second quantization) [16, 213, 72]. In this mode-based picture, Alice controls modes in region A and Bob those in region B, and entanglement reflects correlations across these partitions. Unlike qubits, however, observables from different regions may anticommute, imposing

algebraic constraints that make fermionic entanglement subtler.

In this work we adopt the second-quantization framework and study the extendibility problem for fermionic Gaussian states. We provide a complete characterization of two-sided extendibility for fermionic Gaussian states. Our analysis yields a linear-size semidefinite program, derived directly from their covariance matrices (CM), that gives necessary and sufficient conditions for extendibility. Our framework sharpens separability notions by establishing a consistent operational definition and, in doing so, identifies a distinctive class of *classical–quantum binding states*—including fermionic Gaussian examples—where classical correlations necessarily imply quantum ones. We prove a tight de Finetti-type theorem bounding the trace-norm distance between (k_1, k_2)-extendible and separable states, derive upper bounds on both relative entropy and squashed entanglement, and characterize fermionic Gaussian channels via necessary and sufficient conditions for zero quantum capacity.

Fermionic Gaussian states: We recall the basic framework for fermionic Gaussian states, following standard conventions [26, 29, 200, 78]. An n-mode fermionic system is described by creation and annihilation operators $\{\hat{c}_j^\dagger, \hat{c}_j\}_{j=1}^n$ obeying $\{\hat{c}_j, \hat{c}_k^\dagger\} = \delta_{jk}$, $\{\hat{c}_j, \hat{c}_k\} = 0$. Introducing $2n$ Hermitian Majorana operators $\gamma_{2j-1} = \hat{c}_j + \hat{c}_j^\dagger$, $\gamma_{2j} = i(\hat{c}_j - \hat{c}_j^\dagger)$, they satisfy $\{\gamma_j, \gamma_k\} = 2\delta_{jk}$ and generate the Clifford algebra. The second moments of a state ρ are encoded in its CM M, defined by $M_{qp} = \mathrm{Tr}(i\gamma_q\gamma_p\rho)$ for $q \neq p$ (with $M_{qq} = 0$). A valid CM is (i) real, (ii) antisymmetric, and (iii) satisfies the bona fide condition $I + iM \geq 0$. These conditions are necessary and sufficient for M to be a valid fermionic CM. Fermionic Gaussian states are precisely those fully specified by M. In addition, all physical fermionic states obey the parity superselection rule $[\rho, P] = 0$, with $P = i^n \prod_{j=1}^n \gamma_{2j-1}\gamma_{2j}$ and $P^2 = I$ [16], so the states we consider are physical in this sense.

6.2 Extendibility of fermionic states

Let ρ^{AB} be a (not necessarily Gaussian) bipartite fermionic state on $n = n_A + n_B$ modes. For any physical fermionic state the first moments vanish. Moreover, since Majorana operators are bounded and the Hilbert space is finite-dimensional for finite n, all second moments are finite. The CM then takes the block form

$$M_{AB} = \begin{pmatrix} M_A & X \\ -X^T & M_B \end{pmatrix}, \tag{6.1}$$

where M_A and M_B are the CMs of subsystems A and B, respectively. The requirement that M_{AB} be real and antisymmetric determines the structure of the off-diagonal block:

$X \in \mathbb{R}^{2n_A \times 2n_B}$, which is otherwise arbitrary subject to the usual bona fide condition.

For simplicity we begin by considering k-extendibility on subsystem B. Permutation invariance of $\widetilde{\rho}^{AB_1\ldots B_k}$ constrains the CM to have the block form

$$\widetilde{M}_{AB_1\cdots B_k} = \begin{pmatrix} M_A & X & X & \cdots & X \\ -X^T & M_B & Y & \cdots & Y \\ -X^T & Y & M_B & \cdots & \vdots \\ \vdots & \vdots & \vdots & \ddots & Y \\ -X^T & Y & \cdots & Y & M_B \end{pmatrix}, \tag{6.2}$$

with X real and Y real antisymmetric. Establishing the latter is nontrivial; our argument departs from the bosonic case [122]. As detailed in Appendix E, it exploits the algebra of Majorana operators, and can be adopted for the bosonic setting via field–quadrature relations. Since fermionic Gaussian states are fully determined by their CMs, any bona fide matrix of the form (6.2) represents a k-extendible fermionic Gaussian state. The analogous two-sided CM is given in Appendix E.

Theorem 6.1. *A fermionic Gaussian state ρ_{G}^{AB} is (k_1, k_2)-extendible if and only if it admits a Gaussian (k_1, k_2)-extension.*

Proof. Assume there exists a non-Gaussian extension $\widetilde{\rho}^{AB_1\ldots B_k}$. Its second moments must then satisfy the extension constraints of Eq. 6.2. Since every admissible CM defines a unique fermionic Gaussian state, we may replace $\widetilde{\rho}^{AB_1\ldots B_k}$ by the Gaussian state $\widetilde{\rho}_{\mathrm{G}}^{AB_1\ldots B_k}$ with the same CM. This state is therefore a valid Gaussian extension of ρ_{G}^{AB}. □

An analogous conclusion was reached in the bosonic case using central-limit-theorem techniques [122], and these methods can in fact be adapted to the fermionic setting via a fermionic CLT [50]. Our approach instead uses a direct covariance-matrix argument, yielding a derivation that is more direct and concise without invoking central-limit theorems.

Focusing on the one-sided case, we give a necessary condition for k-extendibility on subsystem B for general fermionic states; the full two-sided result is given in Theorem 6.2. Let ρ^{AB} be a (possibly non-Gaussian) k-extendible state on $n_A + n_B$ modes with CM M_{AB}. Then there exists a CM Δ_B such that

$$I^A \oplus \tfrac{1}{k} I^B + i M_{AB} \;\geq\; \mathbf{0} \oplus i\big(1 - \tfrac{1}{k}\big)\Delta_B. \tag{6.3}$$

If ρ^{AB} is Gaussian, condition (6.3) is also sufficient. We prove this using the block-matrix positivity criterion and its Schur complement [92, Thm. 1.12]:

$$\begin{pmatrix} A & C \\ C^\dagger & B \end{pmatrix} \geq 0 \quad \Leftrightarrow \quad A \geq 0, \quad B \geq C^\dagger A^{-1} C, \tag{6.4}$$

with the singular case handled by standard methods [1].

For a k-extendible state, the CM has the form (6.2) and satisfies $I + i\widetilde{M}_{AB_1\cdots B_k} \geq 0$. Introducing $|+\rangle = \frac{1}{\sqrt{k}} \sum_{j=1}^{k} |j\rangle$, the condition becomes an instance of (6.4) with

$$\begin{aligned} &A = I + iM_A, \;\; C = \sqrt{k}\,(\langle +|_B \otimes iX), \\ &B = k\,|+\rangle\langle +| \otimes iY + I \otimes (I + iM_B - iY). \end{aligned}$$

Since $A \geq 0$, the Schur complement yields

$$\begin{aligned} &(I - |+\rangle\langle +|) \otimes (I + iM_B - iY) \\ &\qquad + |+\rangle\langle +| \otimes \left(I + iM_B + (k-1)iY - kX^T (I + iM_A)^{-1} X\right) \geq 0. \end{aligned}$$

Orthogonality implies each term is positive, giving

$$\frac{1}{k-1}\left[kX^T (I + iM_A)^{-1} X - I - iM_B\right] \leq iY \leq I + iM_B.$$

Defining $\Delta_B := M_B - Y$, the first inequality is equivalent to (6.3), while the second ensures Δ_B is a valid CM.

We now state the full two-sided extension criterion, proved in Appendix E, where we also discuss the landscape of extendibility and show that extendibility from different sides is independent. In addition to providing necessary (and, for Gaussian states, sufficient) conditions, the theorem also enables the explicit construction of extensions as feasible points of the SDP.

Theorem 6.2. *Let ρ^{AB} be a (k_1, k_2)-extendible state (not necessarily Gaussian) with CM (6.1). Then there exist real antisymmetric matrices Δ_A, Δ_B such that $I + i\Delta_A \geq 0$ and $I + i\Delta_B \geq 0$, and*

$$I + i \begin{pmatrix} k_1 M_A - (k_1 - 1)\Delta_A & \sqrt{k_1 k_2}\, X \\ -\sqrt{k_1 k_2}\, X^T & k_2 M_B - (k_2 - 1)\Delta_B \end{pmatrix} \geq 0.$$

If ρ^{AB} is Gaussian, the above semidefinite feasibility condition is both necessary and sufficient.

[1] The case of non-invertible A can be treated by the standard ε-regularization technique [122]; all intermediate reductions and the final result remain valid without assuming invertibility.

Properties of extendible fermionic states.—We now apply Theorem 6.2 to characterize quantum and classical correlations in extendible fermionic Gaussian states. Our first result places a universal constraint on cross–subsystem correlations that applies to *all* fermionic states, not only Gaussians. Its proof is given in Appendix E.

Lemma 6.3. *Let ρ^{AB} be a (k_1, k_2)-extendible state with CM as in Eq. (6.1). The off-diagonal block X of its CM satisfies*

$$X^T X \;\leq\; \frac{4\,I}{k_1 k_2}.$$

This bound progressively suppresses two-point correlations between A and B as the extendibility parameters grow, pushing the CM toward block-diagonal form. In the limit of arbitrary extendibility from either side ($k_1 \to \infty$ or $k_2 \to \infty$), all cross-correlations vanish and the CM factorizes, $M_{AB} = M_A \oplus M_B$, thus providing a proof of the observation made in [213, Observation 1]. From this we deduce:

Proposition 6.4. *A bipartite fermionic Gaussian state ρ_{G}^{AB} is arbitrarily extendible from either side if and only if it is a product state: $\rho_{\mathrm{G}}^{AB} = \rho_{\mathrm{G}}^{A} \otimes \rho_{\mathrm{G}}^{B}$.*

Proof. ($\Rightarrow$) For Gaussian states, Wick's theorem reduces all higher-order correlators to two-point functions. The bound above forces the off-diagonal block X to vanish in the infinite-extendibility limit, eliminating every cross-subsystem correlator and hence yielding $\rho_{\mathrm{G}}^{AB} = \rho_{\mathrm{G}}^{A} \otimes \rho_{\mathrm{G}}^{B}$. ($\Leftarrow$) Conversely, any product Gaussian state is trivially (k_1, k_2)-extendible for all k_1, k_2 by tensoring the marginals. Thus both quantum and classical correlations vanish precisely in the arbitrarily extendible regime. □

This behavior contrasts sharply with bosonic or spin systems [48], where infinite extendibility guarantees separability but not necessarily a product structure. The difference stems from fermionic superselection rules, which restrict admissible decompositions of separable states.

Separable states are convex mixtures of product states, $\rho^{AB} = \sum_x p(x)\, \rho_x^A \otimes \rho_x^B$, but superselection leads to inequivalent notions [16]. Two classes are most relevant: *(i)* in $\mathcal{S}2_\pi$ each component is physical, $\forall x : [\rho_x^A \otimes \rho_x^B, P^{AB}] = 0$, equivalently $[\rho_x^A, P^A] = [\rho_x^B, P^B] = 0$; *(ii)* in $\mathcal{S}2'_\pi$ only the mixture is required to be physical, $[\rho^{AB}, P^{AB}] = 0$, even if the summands are not. Ref. [16] proves the strict inclusion $\mathcal{S}2_\pi \subset \mathcal{S}2'_\pi$.

Guided by our extendibility bounds, we single out $\mathcal{S}2_\pi$ as the operationally natural separability notion for fermions. The key structural fact for Gaussians is the classical–quantum binding property, a term we introduce to denote the absence of purely classical correlations:

Theorem 6.5 (Fermionic Gaussians are classical–quantum binding)**.** *For any fermionic Gaussian state* ρ_G^{AB}*,*

$$\rho_G^{AB} \in \mathcal{S}2_\pi \iff \rho_G^{AB} = \rho_G^A \otimes \rho_G^B.$$

The implication "product $\Rightarrow \mathcal{S}2_\pi$" is immediate. For the converse, we show in Appendix E that parity conservation $[P^A, \rho_{\mathrm{G}}^A] = [P^B, \rho_{\mathrm{G}}^B] = 0$, together with the fact that Majorana operators anticommute with the local parity, $\{P^A, \gamma_m^A\} = \{P^B, \gamma_n^B\} = 0$, implies that $\mathrm{Tr}\big(\gamma_m^A \gamma_n^B \rho_{\mathrm{G}}^{AB}\big) = 0$ for all m, n. Hence the CM is block-diagonal, and in the Gaussian setting such a block structure enforces factorization, $\rho_{\mathrm{G}}^{AB} = \rho_{\mathrm{G}}^A \otimes \rho_{\mathrm{G}}^B$. Together with Proposition 6.4, we conclude that the following statements are equivalent for ρ_{G}^{AB}: (i) $\rho_{\mathrm{G}}^{AB} = \rho_{\mathrm{G}}^A \otimes \rho_{\mathrm{G}}^B$; (ii) ρ_{G}^{AB} is arbitrarily extendible; and (iii) $\rho_{\mathrm{G}}^{AB} \in \mathcal{S}2_\pi$.

Thus, for Gaussian states the separability hierarchy collapses: the $\mathcal{S}2_\pi$ notion coincides with product structure and with arbitrary extendibility. Under weaker notions such as $\mathcal{S}2'_\pi$ or $\mathcal{S}1_\pi$ [16], states may retain residual classical correlations, thereby failing to be extendible; this mismatch makes such notions operationally irrelevant in practice. Finally, while Proposition 6.4 is specific to Gaussians, any non-Gaussian state in $\mathcal{S}2_\pi$ still has vanishing cross second moments (each component is parity-preserving), whereas states in $\mathcal{S}2'_\pi$ can sustain nonzero cross second moments—underscoring $\mathcal{S}2_\pi$ as the operationally justified notion of fermionic separability. In what follows, unless stated otherwise, we use "separable" to mean separable in the $\mathcal{S}2_\pi$ sense, the strongest notion of fermionic separability [16].

Extendible states provide a powerful approximation to the separable set and underlie efficient entanglement criteria [219, 161, 25]. Their importance stems from the quantum de Finetti theorem [183]. For finite k, the central question is how well extendibility approximates separability. Bounding this gap offers an operational measure of the effectiveness of extendibility. Since separability testing and entanglement detection underpin many quantum protocols, our finite de Finetti bounds have direct implications for security analysis and resource estimation in such applications [183, 48]. We now establish finite de Finetti bounds for fermionic Gaussian states, focusing on the two-sided, (k_1, k_2)-extendible case. The following theorem applies uniformly to all separability notions.

Theorem 6.6 (Finite de Finetti bound for fermionic Gaussian states)**.** *Let* ρ_G^{AB} *be a* (k_1, k_2)*-extendible fermionic Gaussian state on* $n_A + n_B$ *modes.*

Define $T := \frac{2}{\sqrt{k_1 k_2}} \min(n_A, n_B, \sqrt{k_1 k_2})$ *and* $h(x) := -x \log_2 x - (1-x)\log_2(1-x)$. *Then*

$$\textit{(i)}\ \|\rho_G^{AB} - \mathrm{SEP}(A\!:\!B)\|_1 \leq T, \tag{6.5}$$

$$\textit{(ii)}\ E_R(\rho_G^{AB}) \leq \tfrac{1}{2}(n_A + n_B)T + h(T/2), \tag{6.6}$$

$$\textit{(iii)}\ E_{\mathrm{sq}}(\rho_G^{AB}) \leq \tfrac{1}{4}(n_A + n_B)T + \tfrac{1}{2}h(T/2). \tag{6.7}$$

Here $\|\rho - \mathrm{SEP}\|_1 := \min_{\sigma \in \mathrm{SEP}(A:B)} \|\rho - \sigma\|_1$ *(trace distance to separable states),* $E_R(\rho)$ *(relative entropy of entanglement)* $:= \inf_{\sigma \in \mathrm{SEP}} S(\rho\|\sigma)$ *with* $S(\rho\|\sigma) = \mathrm{Tr}[\rho(\log\rho - \log\sigma)]$, $E_{\mathrm{sq}}(\rho^{AB})$ *(squashed entanglement)* $:= \frac{1}{2}\inf_{\rho^{ABE}} I(A:B|E)$ *with* $I(A:B|E) = S(AE) + S(BE) - S(E) - S(ABE)$, *and* $S(\tau) = -\mathrm{Tr}[\tau\log\tau]$ *is the von Neumann entropy.*

Proofs are given in Appendix E. These bounds differ sharply from the bosonic and spin cases. For qubits [48] and for general fermionic settings [118, Corollary 15], the distance between extendible and separable states scales exponentially with the number of subsystems or modes. In contrast, our bounds scale linearly in trace norm and only polynomially for E_R and E_{sq}, replacing the previously exponential scaling [118] with a linear one. The bosonic analogue [122, Theorem 3] exploits that any bosonic Gaussian state admits a separable counterpart with a CM proportional to its own [121], ensuring commutativity. No such reduction exists for fermions; instead, our analysis exploits structural properties unique to fermionic CMs. A striking consequence is that, while in the bosonic case the distance in the one-sided k-extendible setting asymptotically vanishes as $1/k$, in the fermionic case it decays only as $1/\sqrt{k}$. We establish the tightness of this scaling with a two-mode example, yielding a lower bound of $1/\sqrt{k_1 k_2}$ and an upper bound of $2/\sqrt{k_1 k_2}$.

Although Theorem 6.6 is stated as a worst-case bound and does not assume any geometric structure, its effective content can be substantially sharper in local settings. If the Gaussian state is the ground state of a local gapped quadratic Hamiltonian, then two-point correlations typically decay exponentially with distance, so the off-diagonal covariance block X across a spatial bipartition is concentrated near the boundary between the two regions. Consequently, the relevant contribution to $\|X\|_1$ comes only from modes within a correlation length of the cut, so the effective number of modes entering the bound is governed by the boundary size $|\partial A| \sim n^{1-1/D}$, rather than the total number of bulk modes $|A| \sim n$. We further study this picture numerically in Appendix E in a geometrically local setting, namely a disordered Kitaev chain, and find that increasing disorder makes the state more likely to be (k_1, k_2)-extendible.

6.3 Extendibility of Fermionic Gaussian Channels

We now turn to fermionic Gaussian channels (FGCs), i.e., completely positive, trace-preserving maps that preserve Gaussianity. As shown in [26], an $m \to n$ fermionic Gaussian channel (FGC) $\mathcal{N}^{A\to B}$ is specified by a real $2n \times 2m$ matrix X and a real antisymmetric $2n \times 2n$ matrix N_B obeying $I + iN_B - XX^T \geq 0$, which implies $XX^T \leq I$ and that N_B is a valid CM. Its action on a CM is $M_A \mapsto XM_AX^T + N_B$. The Choi–Jamiołkowski state $C_{\mathcal{N}}^{AB}$, obtained by applying $\mathcal{N}$ to (the first) half of a maximally entangled fermionic Gaussian state, is itself Gaussian with CM

$$M_{\mathcal{N}} = \begin{pmatrix} N_B & X \\ -X^T & 0 \end{pmatrix}. \tag{6.8}$$

Determining when a quantum channel has zero capacity is a fundamental problem in quantum information. A prominent example is the class of antidegradable channels, where the environment can reconstruct the channel output; by the quantum no-cloning theorem, their quantum capacity vanishes. Characterizing antidegradability has thus attracted significant interest [54, 108, 153, 78].

Antidegradability is closely related to k-extendible channels: a channel $\mathcal{N}$ is k-extendible if its Choi state $C_{\mathcal{N}}^{AB}$ is k-extendible with respect to the output system B [106, 105]. In particular, antidegradability is equivalent to 2-extendibility of the Choi state [158]. Applying Theorem 6.2 to the Choi CM $M_{\mathcal{N}}$ in Eq. (6.8) with $k_1 = 2$, $k_2 = 1$, $M_B = 0$, and $M_A = N_B$ yields a necessary and sufficient condition for antidegradability of fermionic Gaussian channels:

Proposition 6.7. *A fermionic Gaussian channel (X, N_B) is antidegradable if and only if there exists a real antisymmetric matrix Δ satisfying*

$$i\Delta \;\leq\; \left\{ I,\, I + 2iN_B - 2XX^T \right\}.$$

A simple example illustrating Proposition 6.7 is the fermionic Gaussian attenuation channel. Take $X = \sqrt{\lambda}\,\mathbb{I}$ and $N_B = 0$, for example on a two-mode system. The channel acts on covariance matrices as $\Gamma_{\text{out}} = \lambda\,\Gamma_{\text{in}}$, and the complete-positivity condition holds exactly for $0 \leq \lambda \leq 1$. For these parameters, the matrix appearing in Proposition 6.7 is $I + 2iN_B - 2XX^{\mathsf{T}} = I - 2\lambda I = (1 - 2\lambda)I$, from which one finds that the channel is antidegradable precisely when $0 \leq \lambda \leq \frac{1}{2}$. Thus sufficiently strong attenuation places the channel in the antidegradable regime, where its quantum capacity is zero. From an experimental perspective, the conditions entering these results depend only on the channel

action on second moments, and the required matrices X and N_B can therefore be inferred from covariance-matrix tomography.

Two other important zero-capacity classes are *replacement channels*, which map every input to a fixed Gaussian state, and *entanglement-breaking* (EB) channels, which destroy all input–output entanglement. A channel is EB if and only if its Choi state is separable [93]. In the fermionic Gaussian setting, Gaussian separability is extremely restrictive: any "measure-and-prepare" map that outputs Gaussian states must prepare the *same* Gaussian state for all measurement outcomes. Intuitively, this follows because fermionic Gaussian states have vanishing first moments (due to parity superselection) and mixtures of Gaussian states with distinct covariances are generically non-Gaussian. Consequently, the only Gaussian-preserving EB maps have $X = 0$, so that $M_A \mapsto N_B$ is a replacement channel. We now give a proof using the extendibility framework.

Proposition 6.8. *For fermionic Gaussian channels, entanglement-breaking is equivalent to being a replacement channel.*

Proof. Applying Theorem 6.2 to the Choi CM $M_{\mathcal{N}}$ in Eq. (6.8) with $k_1 = k$, $k_2 = 1$, $M_B = 0$, and $M_A = N_B$, and using the Schur complement with respect to the bottom-right block gives $I + iN_B + i(k-1)Z - kXX^T \geq 0$. Taking $k \to \infty$ yields $XX^T \leq iZ$; combined with $XX^T \leq -iZ$, we conclude $X = 0$. The channel then reduces to $M_A \mapsto N_B$, a replacement channel, which is trivially entanglement-breaking. □

The coincidence of EB and replacement channels may have interesting consequences in quantum information processing. In particular, it implies that, asymptotically, discriminating any (possibly non-Gaussian) fermionic channel from a fermionic Gaussian EB channel requires only non-adaptive tensor-power strategies [52].

6.4 Discussion

We have presented a comprehensive study of two-sided (k_1, k_2)-extendibility for fermionic Gaussian states, providing a full covariance-matrix characterization and a simple SDP formulation. Leveraging tools from quantum information theory, we derived finite de Finetti bounds that improve previous exponential scaling to linear in the number of modes, with complementary bounds for relative entropy and squashed entanglement. These results provide operational criteria for separability in the Gaussian sector, ruling out nontrivial separable Gaussian states and identifying the definition that is both natural and practically

meaningful. Finally, we established an SDP criterion for the anti-degradability of fermionic Gaussian channels.

Promising directions for future work include exploring applications of the derived bounds, investigating the family of classical–quantum binding states, and performing a systematic study of two-sided (k_1, k_2)-extendibility, which remains largely unexplored. Finally, it is an open question whether fermionic Gaussian steerability can be rigorously defined and quantified, and how it is constrained by finite extendibility [122, 116].

References

[1] Scott Aaronson and Daniel Gottesman. Improved simulation of stabilizer circuits. *Physical Review A*, 70(5):052328, 2004.

[2] Dorit Aharonov. Quantum to classical phase transition in noisy quantum computers. *Phys. Rev. A*, 62:062311, Nov 2000.

[3] Dorit Aharonov and Michael Ben-Or. Fault-tolerant quantum computation with constant error rate, 1999.

[4] Google Quantum AI and collaborators. Quantum error correction below the surface code threshold, 2024.

[5] Asma Al-Qasimi and Daniel F. V. James. Sudden death of entanglement at finite temperature. *Phys. Rev. A*, 77:012117, Jan 2008.

[6] Victor V. Albert. Lindbladians with multiple steady states: theory and applications, January 2018. arXiv:1802.00010 [cond-mat, physics:math-ph, physics:quant-ph].

[7] Victor V. Albert and Liang Jiang. Symmetries and conserved quantities in lindblad master equations. *Physical Review A*, 89(2), February 2014.

[8] Luigi Amico, Rosario Fazio, Andreas Osterloh, and Vlatko Vedral. Entanglement in many-body systems. *Rev. Mod. Phys.*, 80:517–576, May 2008.

[9] Leandro Aolita, Fernando de Melo, and Luiz Davidovich. Open-system dynamics of entanglement:a key issues review. *Reports on Progress in Physics*, 78(4):042001, March 2015.

[10] M. C. Arnesen, S. Bose, and V. Vedral. Natural thermal and magnetic entanglement in the 1d heisenberg model. *Physical Review Letters*, 87(1), June 2001.

[11] Koenraad M R Audenaert. A sharp continuity estimate for the von neumann entropy. *Journal of Physics A: Mathematical and Theoretical*, 40(28):8127, jun 2007.

[12] Ainesh Bakshi, Allen Liu, Ankur Moitra, and Ewin Tang. High-temperature gibbs states are unentangled and efficiently preparable, 2025.

[13] Yimu Bao, Soonwon Choi, and Ehud Altman. Theory of the phase transition in random unitary circuits with measurements. *Phys. Rev. B*, 101:104301, Mar 2020.

[14] Yimu Bao, Ruihua Fan, Ashvin Vishwanath, and Ehud Altman. Mixed-state topological order and the errorfield double formulation of decoherence-induced transitions, 2023.

[15] Stephen D. Bartlett, Terry Rudolph, and Robert W. Spekkens. Reference frames, superselection rules, and quantum information. *Reviews of Modern Physics*, 79(2):555–609, April 2007. Publisher: American Physical Society.

[16] Mari-Carmen Bañuls, J. Ignacio Cirac, and Michael M. Wolf. Entanglement in fermionic systems. *Physical Review A*, 76(2), August 2007.

[17] Jan Behrends, Florian Venn, and Benjamin Béri. Surface codes, quantum circuits, and entanglement phases, 2022.

[18] F. Benatti, R. Floreanini, and U. Marzolino. Entanglement in fermion systems and quantum metrology. *Physical Review A*, 89(3), March 2014.

[19] Ingemar Bengtsson and Karol Życzkowski. *Geometry of Quantum States: An Introduction to Quantum Entanglement*. Cambridge University Press, August 2017.

[20] Lennart Bittel, Antonio Anna Mele, Jens Eisert, and Lorenzo Leone. Optimal trace-distance bounds for free-fermionic states: Testing and improved tomography, 2025.

[21] A. Bolt, G. Duclos-Cianci, D. Poulin, and T.M. Stace. Foliated quantum error-correcting codes. *Physical Review Letters*, 117(7), August 2016.

[22] H. Bombin and M. A. Martin-Delgado. Optimal resources for topological two-dimensional stabilizer codes: Comparative study. *Phys. Rev. A*, 76:012305, Jul 2007.

[23] Hector Bombin, Daniel Litinski, Naomi Nickerson, Fernando Pastawski, and Sam Roberts. Unifying flavors of fault tolerance with the ZX calculus. *Quantum*, 8:1379, June 2024.

[24] Thomas Botzung, M. Buchhold, S. Diehl, and Markus Müller. Robustness and measurement-induced percolation of the surface code. *J. Phys. A: Math. Theor.*, 58(20):205304, 2025.

[25] Fernando G.S.L. Brandão, Matthias Christandl, and Jon Yard. A quasipolynomial-time algorithm for the quantum separability problem. In *Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing*, STOC '11, page 343–352, New York, NY, USA, 2011. Association for Computing Machinery.

[26] Sergey Bravyi. Lagrangian representation for fermionic linear optics, 2004.

[27] Sergey Bravyi. Classical capacity of fermionic product channels, 2005.

[28] Sergey Bravyi, Matthias Englbrecht, Robert König, and Nolan Peard. Correcting coherent errors with surface codes. *npj Quantum Inf.*, 4:55, 10 2018.

[29] Sergey Bravyi and David Gosset. Complexity of quantum impurity problems. *Communications in Mathematical Physics*, 356(2):451–500, December 2017.

[30] Sergey Bravyi, David Gosset, Robert König, and Kristan Temme. Approximation algorithms for quantum many-body problems. *Journal of Mathematical Physics*, 60(3):032203, 03 2019.

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

[32] Sergey Bravyi and Robert König. Classification of topologically protected gates for local stabilizer codes. *Phys. Rev. Lett.*, 110(17), April 2013.

[33] Sergey Bravyi and Robert Raussendorf. Measurement-based quantum computation with the toric code states. *Physical Review A*, 76(2), aug 2007.

[34] Hans J. Briegel and Robert Raussendorf. Persistent Entanglement in Arrays of Interacting Particles. *Physical Review Letters*, 86(5):910–913, January 2001. Publisher: American Physical Society.

[35] Benjamin J. Brown and Sam Roberts. Universal fault-tolerant measurement-based quantum computation. *Physical Review Research*, 2(3), August 2020.

[36] Winton Brown and David Poulin. Quantum markov networks and commuting hamiltonians, 2012.

[37] Nicolas Brunner and Tamás Vértesi. Persistency of entanglement and nonlocality in multipartite quantum systems. *Physical Review A*, 86(4), October 2012.

[38] Berislav Buča and Tomaž Prosen. A note on symmetry reductions of the lindblad equation: transport in constrained open spin chains. *New Journal of Physics*, 14(7):073007, July 2012.

[39] Carlton M. Caves, Christopher A. Fuchs, and Pranaw Rungta. Entanglement of Formation of an Arbitrary State of Two Rebits. *Foundations of Physics Letters*, 14(3):199–212, June 2001.

[40] Daniel Centeno, Marco Erba, Thomas D. Galley, David Schmid, John H. Selby, Robert W. Spekkens, Sina Soltani, Jacopo Surace, Alex Wilce, and Yìlè Yīng. Twirled worlds: symmetry-induced failures of tomographic locality, July 2025. arXiv:2407.21688 [quant-ph] version: 3.

[41] Chi-Fang Chen, Michael J. Kastoryano, Fernando G. S. L. Brandão, and András Gilyén. Quantum thermal state preparation, 2023.

[42] Chi-Fang Chen, Michael J. Kastoryano, and András Gilyén. An efficient and exact noncommutative quantum gibbs sampler, 2023.

[43] Chi-Fang Chen and Cambyse Rouzé. Quantum gibbs states are locally markovian, 2025.

[44] Yu-Hsueh Chen and Tarun Grover. Separability transitions in topological states induced by local decoherence. *arXiv preprint arXiv:2309.11879*, 2023.

[45] Yu-Hsueh Chen and Tarun Grover. Symmetry-enforced many-body separability transitions. *PRX Quantum*, 5:030310, Jul 2024.

[46] Yu-Hsueh Chen and Tarun Grover. Unconventional topological mixed-state transition and critical phase induced by self-dual coherent errors, 2024.

[47] Wonjune Choi, Michael Knap, and Frank Pollmann. Finite-temperature entanglement negativity of fermionic symmetry-protected topological phases and quantum critical points in one dimension. *Physical Review B*, 109(11), March 2024.

[48] Matthias Christandl, Robert König, Graeme Mitchison, and Renato Renner. One-and-a-Half quantum de finetti theorems. *Communications in Mathematical Physics*, 273(2):473–498, July 2007.

[49] Christopher T. Chubb and Steven T. Flammia. Statistical mechanical models for quantum codes with correlated noise. *Annales de L'Institut Henri Poincaré D*, 8:269–321, 2021.

[50] Luke Coffman, Graeme Smith, and Xun Gao. Measuring non-gaussian magic in fermions: Convolution, entropy, and the violation of wick's theorem and the matchgate identity, 2025.

[51] Valerie Coffman, Joydip Kundu, and William K. Wootters. Distributed entanglement. *Phys. Rev. A*, 61:052306, Apr 2000.

[52] Tom Cooney, Milán Mosonyi, and Mark M. Wilde. Strong converse exponents for a quantum channel discrimination problem and quantum-feedback-assisted communication. *Communications in Mathematical Physics*, 344(3):797–829, May 2016.

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

[54] Toby S. Cubitt, Mary Beth Ruskai, and Graeme Smith. The structure of degradable quantum channels. *Journal of Mathematical Physics*, 49(10):102104, 10 2008.

[55] Caroline de Groot, Alex Turzillo, and Norbert Schuch. Symmetry protected topological order in open quantum systems. *Quantum*, 6:856, 2022.

[56] Nicolas Delfosse and Adam Paetznick. Spacetime codes of clifford circuits, 2023.

[57] Eric Dennis, Alexei Kitaev, Andrew Landahl, and John Preskill. Topological quantum memory. *Journal of Mathematical Physics*, 43(9):4452–4505, September 2002.

[58] S. Diehl, A. Micheli, A. Kantian, B. Kraus, H. P. Büchler, and P. Zoller. Quantum states and phases in driven open quantum systems with cold atoms. *Nature Physics*, 4(11):878–883, 2008.

[59] A. C. Doherty, Pablo A. Parrilo, and Federico M. Spedalieri. Distinguishing separable and entangled states. *Phys. Rev. Lett.*, 88:187904, Apr 2002.

[60] Andrew C. Doherty, Pablo A. Parrilo, and Federico M. Spedalieri. Complete family of separability criteria. *Phys. Rev. A*, 69:022308, Feb 2004.

[61] Andrew C. Doherty, Pablo A. Parrilo, and Federico M. Spedalieri. Detecting multipartite entanglement. *Phys. Rev. A*, 71:032333, Mar 2005.

[62] Jonathan P. Dowling and Gerard J. Milburn. Quantum technology: the second quantum revolution. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 361(1809):1655–1674, 2003.

[63] Finn Eckstein, Bo Han, Simon Trebst, and Guo-Yi Zhu. Learning transitions of topological surface codes, 2025.

[64] J. Eisert, M. Cramer, and M. B. Plenio. Colloquium: Area laws for the entanglement entropy. *Rev. Mod. Phys.*, 82:277–306, Feb 2010.

[65] Jens Eisert, Viktor Eisler, and Zoltán Zimborás. Entanglement negativity bounds for fermionic gaussian states. *Physical Review B*, 97(16), April 2018.

[66] Viktor Eisler and Zoltán Zimborás. On the partial transpose of fermionic gaussian states. *New Journal of Physics*, 17(5):053048, May 2015.

[67] T. D. Ellison and M. Cheng. Toward a classification of mixed-state topological orders in two dimensions. *PRX Quantum*, 6:010315, 2025.

[68] Ruihua Fan, Yimu Bao, Ehud Altman, and Ashvin Vishwanath. Diagnostics of mixed-state topological order and breakdown of quantum memory. *PRX Quantum*, 5:020343, May 2024.

[69] M. Fannes, J. T. Lewis, and A. Verbeure. Symmetric states of composite systems. *Letters in Mathematical Physics*, 15(3):255–260, April 1988.

[70] B. V. Fine, F. Mintert, and A. Buchleitner. Equilibrium entanglement vanishes at finite temperature. *Physical Review B*, 71(15), April 2005.

[71] Matthew P.A. Fisher, Vedika Khemani, Adam Nahum, and Sagar Vijay. Random quantum circuits. *Annual Review of Condensed Matter Physics*, 14(1):335–379, mar 2023.

[72] Nicolai Friis, Antony R. Lee, and David Edward Bruschi. Fermionic-mode entanglement in quantum information. *Phys. Rev. A*, 87:022338, Feb 2013.

[73] Samuel J. Garratt and Max McGinley. Entanglement and private information in many-body thermal states, 2025.

[74] Domenico Giulini. Superselection Rules, May 2009. arXiv:0710.1516 [quant-ph].

[75] Daniel Gottesman. The heisenberg representation of quantum computers. *arXiv preprint quant-ph/9807006*, 1998.

[76] Daniel Gottesman. Opportunities and challenges in fault-tolerant quantum computation, 2022.

[77] M. Grassl, T. Beth, and T. Pellizzari. Codes for the quantum erasure channel. *Phys. Rev. A*, 56:33, 1997.

[78] Eliška Greplová and Géza Giedke. Degradability of fermionic gaussian channels. *Physical Review Letters*, 121(20), November 2018.

[79] Michael J. Gullans and David A. Huse. Dynamical purification phase transition induced by quantum measurements. *Phys. Rev. X*, 10:041020, Oct 2020.

[80] Yuchen Guo, Jian-Hao Zhang, Zhen Bi, and Shuo Yang. Triggering boundary phase transitions through bulk measurements in 2d cluster states, 2023.

[81] Yuchen Guo, Jian-Hao Zhang, Shuo Yang, and Zhen Bi. Locally purified density operators for symmetry-protected topological phases in mixed states. *arXiv preprint arXiv:2403.16978*, 2024.

[82] Leonid Gurvits and Howard Barnum. Separable balls around the maximally mixed multipartite quantum states. *Physical Review A*, 68(4):042312, October 2003.

[83] Gábor B. Halász, Timothy H. Hsieh, and Leon Balents. Fracton topological phases from strongly coupled spin chains. *Phys. Rev. Lett.*, 119:257202, Dec 2017.

[84] O. Hart and C. Castelnovo. Entanglement negativity and sudden death in the toric code at finite temperature. *Physical Review B*, 97(14), April 2018.

[85] Matthew B Hastings. Topological order at nonzero temperature. *Physical review letters*, 107(21):210501, 2011.

[86] Matthew B. Hastings and Jeongwan Haah. Dynamically Generated Logical Qubits. *Quantum*, 5:564, October 2021.

[87] Matthew B. Hastings and Tohru Koma. Spectral gap and exponential decay of correlations. *Commun. Math. Phys.*, 265(3):781–804, April 2006.

[88] Jacob Hauser, Yimu Bao, Shengqi Sang, Ali Lavasani, Utkarsh Agrawal, and Matthew P. A. Fisher. Information dynamics in decohered quantum memory with repeated syndrome measurements: a dual approach, 2024.

[89] Patrick Hayden, Debbie W Leung, and Andreas Winter. Aspects of generic entanglement. *Communications in Mathematical Physics*, 265(1):95–117, July 2006.

[90] A. J. Heeger, S. Kivelson, J. R. Schrieffer, and W. P. Su. Solitons in conducting polymers. *Rev. Mod. Phys.*, 60:781–850, Jul 1988.

[91] G. C. Hegerfeldt, K. Kraus, and E. P. Wigner. Proof of the Fermion Superselection Rule without the Assumption of Time-Reversal Invariance. *Journal of Mathematical Physics*, 9(12):2029–2031, December 1968.

[92] Roger A. Horn and Fuzhen Zhang. *Basic Properties of the Schur Complement*, pages 17–46. Springer US, Boston, MA, 2005.

[93] Michael Horodecki, Peter W. Shor, and Mary Beth Ruskai. Entanglement breaking channels. *Reviews in Mathematical Physics*, 15(06):629–641, 2003.

[94] Michał Horodecki, Paweł Horodecki, and Ryszard Horodecki. Separability of mixed states: necessary and sufficient conditions. *Physics Letters A*, 223(1–2):1–8, November 1996.

[95] Michał Horodecki, Jonathan Oppenheim, and Andreas Winter. Partial quantum information. *Nature*, 436(7051):673–676, August 2005.

[96] Ryszard Horodecki, Paweł Horodecki, Michał Horodecki, and Karol Horodecki. Quantum entanglement. *Rev. Mod. Phys.*, 81:865–942, Jun 2009.

[97] Tim Hsieh. Kitp tutorial "mixed-state phases", 2025.

[98] Timothy H. Hsieh and Gábor B. Halász. Fractons from partons. *Phys. Rev. B*, 96:165105, Oct 2017.

[99] Po-Shen Hsin, Ryohei Kobayashi, and Abhinav Prem. Higher-Form Anomalies Imply Intrinsic Long-Range Entanglement, April 2025. arXiv:2504.10569 [quant-ph].

[100] Ze-Min Huang, Luis Colmenarez, Markus Müller, and Sebastian Diehl. Coherent information as a mixed-state topological order parameter of fermions. *Physical Review Research*, 7(4), October 2025.

[101] Chao-Ming Jian, Bela Bauer, Anna Keselman, and Andreas W. W. Ludwig. Criticality and entanglement in non-unitary quantum circuits and tensor networks of non-interacting fermions, 2022.

[102] Chao-Ming Jian, Hassan Shapourian, Bela Bauer, and Andreas W. W. Ludwig. Measurement-induced entanglement transitions in quantum circuits of non-interacting fermions: Born-rule versus forced measurements, 2023.

[103] Chao-Ming Jian, Yi-Zhuang You, Romain Vasseur, and Andreas W. W. Ludwig. Measurement-induced criticality in random quantum circuits. *Phys. Rev. B*, 101:104302, Mar 2020.

[104] Kohtaro Kato and Fernando G. S. L. Brandão. Quantum approximate markov chains are thermal. *Communications in Mathematical Physics*, 370(1):117–149, June 2019.

[105] Eneet Kaur, Siddhartha Das, Mark M. Wilde, and Andreas Winter. Extendibility limits the performance of quantum processors. *Phys. Rev. Lett.*, 123:070502, Aug 2019.

[106] Eneet Kaur, Siddhartha Das, Mark M. Wilde, and Andreas Winter. Resource theory of unextendibility and nonasymptotic quantum capacity. *Phys. Rev. A*, 104:022401, Aug 2021.

[107] Markus S. Kesselring, Julio C. Magdalena de la Fuente, Felix Thomsen, Jens Eisert, Stephen D. Bartlett, and Benjamin J. Brown. Anyon condensation and the color code. *PRX Quantum*, 5:010342, Mar 2024.

[108] Zahra Baghali Khanian and Christoph Hirche. On strong converse bounds for the private and quantum capacities of anti-degradable channels, 2025.

[109] Yonna Kim, Ali Lavasani, and Sagar Vijay. Persistent topological negativity in a high-temperature mixed-state, 2025.

[110] Alexei Kitaev. *Fault-tolerant quantum computation by anyons*. 1997.

[111] Alexei Kitaev. Anyons in an exactly solved model and beyond. *Annals of Physics*, 321(1):2–111, jan 2006.

[112] A.Yu. Kitaev. Fault-tolerant quantum computation by anyons. *Annals of Physics*, 303(1):2–30, January 2003.

[113] E. Knill. Scalable quantum computing in the presence of large detected-error rates. *Phys. Rev. A*, 71:042322, 2005.

[114] Emanuel Knill, Raymond Laflamme, and Wojciech H. Zurek. Resilient quantum computation: error models and thresholds. *Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*, 454(1969):365–384, January 1998.

[115] Masato Koashi and Andreas Winter. Monogamy of quantum entanglement and other correlations. *Phys. Rev. A*, 69:022309, Feb 2004.

[116] Ioannis Kogias, Antony R. Lee, Sammy Ragy, and Gerardo Adesso. Quantification of gaussian quantum steering. *Phys. Rev. Lett.*, 114:060403, Feb 2015.

[117] Christian Krumnow, Zoltán Zimborás, and Jens Eisert. A fermionic de finetti theorem. *Journal of Mathematical Physics*, 58(12), December 2017.

[118] Christian Krumnow, Zoltán Zimborás, and Jens Eisert. Extendibility of fermionic states and rigorous ground state approximations of interacting fermionic systems, 2024.

[119] Tomotaka Kuwahara. Clustering of conditional mutual information and quantum markov structure at arbitrary temperatures, 2024.

[120] Tomotaka Kuwahara, Álvaro M. Alhambra, and Anurag Anshu. Improved thermal area law and quasilinear time algorithm for quantum gibbs states. *Physical Review X*, 11(1), March 2021.

[121] Robert König and Michael M. Wolf. On exchangeable continuous variable systems. *Journal of Mathematical Physics*, 50(1):012102, 01 2009.

[122] Ludovico Lami, Sumeet Khatri, Gerardo Adesso, and Mark M. Wilde. Extendibility of bosonic gaussian states. *Physical Review Letters*, 123(5), July 2019.

[123] Cécilia Lancien, Sara Di Martino, Marcus Huber, Marco Piani, Gerardo Adesso, and Andreas Winter. Should entanglement measures be monogamous or faithful? *Phys. Rev. Lett.*, 117:060501, Aug 2016.

[124] Walter Langel. T. helgaker, p. jorgensen, j. olsen, t. helgaker: Molecular electronic-structure theory. *Journal of Solid State Electrochemistry*, 6(5):359–360, June 2002.

[125] Ali Lavasani, Yahya Alavirad, and Maissam Barkeshli. Measurement-induced topological entanglement transitions in symmetric random quantum circuits. *Nature Phys.*, 17:342–347, 2021.

[126] Ali Lavasani, Zhu-Xi Luo, and Sagar Vijay. Monitored quantum dynamics and the kitaev spin liquid. *arXiv preprint arXiv:2207.02877*, 2022.

[127] Dongjin Lee and Beni Yoshida. Randomly monitored quantum codes, 2024.

[128] Jong Yeon Lee. Exact calculations of coherent information for toric codes under decoherence: Identifying the fundamental error threshold. *Physical Review Letters*, 134(25), June 2025.

[129] Jong Yeon Lee, Wenjie Ji, Zhen Bi, and Matthew P. A. Fisher. Decoding measurement-prepared quantum phases and transitions: from ising model to gauge theory, and beyond, 2022.

[130] Jong Yeon Lee, Chao-Ming Jian, and Cenke Xu. Quantum criticality under decoherence or weak measurement. *arXiv preprint arXiv:2301.05238*, 2023.

[131] Jong Yeon Lee, Yi-Zhuang You, and Cenke Xu. Symmetry protected topological phases under decoherence. *arXiv preprint arXiv:2210.16323*, 2022.

[132] Leonardo A. Lessa, Meng Cheng, and Chong Wang. Mixed-state quantum anomaly and multipartite entanglement, 2024.

[133] Leonardo A. Lessa, Ruochen Ma, Jian-Hao Zhang, Zhen Bi, Meng Cheng, and Chong Wang. Strong-to-Weak Spontaneous Symmetry Breaking in Mixed Quantum States, November 2024. arXiv:2405.03639 [quant-ph].

[134] Leonardo A. Lessa, Shengqi Sang, Tsung-Cheng Lu, Timothy H. Hsieh, and Chong Wang. Higher-form anomaly and long-range entanglement of mixed states, March 2025. arXiv:2503.12792 [quant-ph].

[135] Péter Lévay, Szilvia Nagy, and János Pipek. Elementary formula for entanglement entropies of fermionic systems. *Phys. Rev. A*, 72:022302, Aug 2005.

[136] Michael Levin and Xiao-Gang Wen. Fermions, strings, and gauge fields in lattice spin models. *Physical Review B*, 67(24), jun 2003.

[137] Kangle Li and Liujun Zou. Symmetry-enforced minimal entanglement and correlation in quantum spin chains, December 2024. arXiv:2412.20765 [cond-mat].

[138] Ke Li and Andreas Winter. Squashed entanglement, k-extendibility, quantum markov chains, and recovery maps. *Foundations of Physics*, 48(8):910–924, August 2018.

[139] Yahui Li, Frank Pollmann, Nicholas Read, and Pablo Sala. Highly entangled stationary states from strong symmetries. *Physical Review X*, 15(1), March 2025.

[140] Yaodong Li, Xiao Chen, and Matthew P. A. Fisher. Quantum zeno effect and the many-body entanglement transition. *Phys. Rev. B*, 98:205136, Nov 2018.

[141] Yaodong Li, Xiao Chen, and Matthew P. A. Fisher. Measurement-driven entanglement transition in hybrid quantum circuits. *Phys. Rev. B*, 100:134306, Oct 2019.

[142] Yaodong Li, Xiao Chen, Andreas W. W. Ludwig, and Matthew P. A. Fisher. Conformal invariance and quantum nonlocality in critical hybrid circuits. *Physical Review B*, 104(10), sep 2021.

[143] Yuan Li and Yu-E. Li. Some characterizations of the trace norm triangle equality. *Linear Algebra and its Applications*, 484:396–408, 2015.

[144] Zhi Li, Dongjin Lee, and Beni Yoshida. How much entanglement is needed for emergent anyons and fermions?, May 2024. arXiv:2405.07970.

[145] Cheng-Ju Lin, Weicheng Ye, Yijian Zou, Shengqi Sang, and Timothy H. Hsieh. Probing sign structure using measurement-induced entanglement. *Quantum*, 7:910, feb 2023.

[146] Hanchen Liu and Xiao Chen. Coherent error induced phase transition, 2025.

[147] Hanchen Liu, Tianci Zhou, and Xiao Chen. Measurement-induced entanglement transition in a two-dimensional shallow circuit. *Phys. Rev. B*, 106:144311, Oct 2022.

[148] Shuangyuan Lu, Penghao Zhu, and Yuan-Ming Lu. Bilayer construction for mixed state phenomena with strong, weak symmetries and symmetry breakings, 2024.

[149] Tsung-Cheng Lu, Zhehao Zhang, Sagar Vijay, and Timothy H. Hsieh. Mixed-state long-range order and criticality from measurement and feedback. *PRX Quantum*, 4:030318, Aug 2023.

[150] Ruochen Ma and Alex Turzillo. Symmetry protected topological phases of mixed states in the doubled space. *arXiv preprint arXiv:2403.13280*, 2024.

[151] Ruochen Ma and Chong Wang. Average symmetry-protected topological phases. *Physical Review X*, 13(3):031016, 2023.

[152] Ruochen Ma, Jian-Hao Zhang, Zhen Bi, Meng Cheng, and Chong Wang. Topological phases with average symmetries: the decohered, the disordered, and the intrinsic, 2023.

[153] Francesco Anna Mele, Farzin Salek, Vittorio Giovannetti, and Ludovico Lami. Quantum communication on the bosonic loss-dephasing channel. *Phys. Rev. A*, 110:012460, Jul 2024.

[154] Joseph Merritt and Lukasz Fidkowski. Entanglement transitions with free fermions. *Physical Review B*, 107(6), feb 2023.

[155] Spyridon Michalakis and Justyna P Zwolak. Stability of frustration-free hamiltonians. *Communications in Mathematical Physics*, 322:277–302, 2013.

[156] Amin Moharramipour, Leonardo A. Lessa, Chong Wang, Timothy H. Hsieh, and Subhayan Sahu. Symmetry-enforced entanglement in maximally mixed states. *PRX Quantum*, 5(4), December 2024.

[157] Hajime Moriya. On separable states for composite systems of distinguishable fermions, 2005.

[158] Geir Ove Myhr and Norbert Lütkenhaus. Spectrum conditions for symmetric extendible states. *Phys. Rev. A*, 79:062307, Jun 2009.

[159] Adam Nahum, Jonathan Ruhman, Sagar Vijay, and Jeongwan Haah. Quantum entanglement growth under random unitary dynamics. *Physical Review X*, 7(3), jul 2017.

[160] Adam Nahum, P. Serna, A. M. Somoza, and M. Ortuñ o. Loop models with crossings. *Physical Review B*, 87(18), may 2013.

[161] Miguel Navascués, Masaki Owari, and Martin B. Plenio. Power of symmetric extensions for entanglement detection. *Phys. Rev. A*, 80:052306, Nov 2009.

[162] Amir-Reza Negari, Tyler D. Ellison, and Timothy H. Hsieh. Spacetime markov length: a diagnostic for fault tolerance via mixed-state phases, 2025.

[163] Amir-Reza Negari, Leonardo A. Lessa, and Subhayan Sahu. Symmetry enforces entanglement at high temperatures, 2025.

[164] Amir-Reza Negari, Subhayan Sahu, Jan Behrends, Benjamin Béri, and Timothy H. Hsieh. Critical non-equilibrium phases from noisy topological memories, 2026.

[165] Amir-Reza Negari, Subhayan Sahu, and Timothy H. Hsieh. Measurement-induced phase transitions in the toric code. *Phys. Rev. B*, 109:125148, Mar 2024.

[166] Amir-Reza Negari and Farzin Salek. Extendibility of fermionic gaussian states, 2025.

[167] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press, 2010.

[168] Takuya Okuda, Aswin Parayil Mana, and Hiroki Sukeno. Anomaly inflow for css and fractonic lattice models and dualities via cluster state measurement, 2024.

[169] Stefano Paesani and Benjamin J. Brown. High-threshold quantum computing by fusing one-dimensional cluster states. *Phys. Rev. Lett.*, 131:120603, Sep 2023.

[170] Gilles Parez and William Witczak-Krempa. The fate of entanglement, 2025.

[171] Asher Peres. Separability criterion for density matrices. *Physical Review Letters*, 77(8):1413–1415, August 1996.

[172] Ingo Peschel. Calculation of reduced density matrices from correlation functions. *Journal of Physics A: Mathematical and General*, 36(14):L205, mar 2003.

[173] Ingo Peschel and Viktor Eisler. Reduced density matrices and entanglement entropy in free lattice models. *Journal of Physics A: Mathematical and Theoretical*, 42(50):504003, dec 2009.

[174] Martin B. Plbnio and Shashank Virmani. An introduction to entanglement measures. *Quantum Info. Comput.*, 7(1):1–51, January 2007.

[175] Andrew C. Potter and Romain Vasseur. *Entanglement Dynamics in Hybrid Quantum Circuits*, pages 211–249. Springer International Publishing, Cham, 2022.

[176] Marvin Qi, Leo Radzihovsky, and Michael Hermele. Fracton phases via exotic higher-form symmetry-breaking. *Annals of Physics*, 424:168360, 2021.

[177] B. V. Rajarama Bhat, K. R. Parthasarathy, and Ritabrata Sengupta. On the equivalence of separability and extendability of quantum states. *Reviews in Mathematical Physics*, 29(04):1750012, 2017.

[178] Tibor Rakovszky, Sarang Gopalakrishnan, and Curt von Keyserlingk. Defining stable phases of open quantum systems. *Phys. Rev. X*, 14:041031, Nov 2024.

[179] Akshar Ramkumar, Yiyi Cai, Yu Tong, and Jiaqing Jiang. High-temperature fermionic gibbs states are mixtures of gaussian states, 2025.

[180] R. Raussendorf, D. E. Browne, and H. J. Briegel. Measurement-based quantum computation with cluster states. *Physical Review A*, 68(2):022312, August 2003. arXiv:quant-ph/0301052.

[181] Robert Raussendorf, Sergey Bravyi, and Jim Harrington. Long-range quantum entanglement in noisy cluster states. *Phys. Rev. A*, 71:062313, Jun 2005.

[182] Robert Raussendorf and Hans J. Briegel. A one-way quantum computer. *Phys. Rev. Lett.*, 86:5188–5191, May 2001.

[183] Renato Renner. Symmetry of large physical systems implies independence of subsystems. *Nature Physics*, 3(9):645–649, September 2007.

[184] Sam Roberts and Stephen D. Bartlett. Symmetry-protected self-correcting quantum memories. *Phys. Rev. X*, 10:031041, Aug 2020.

[185] Cambyse Rouzé, Daniel Stilck França, and Álvaro M. Alhambra. Efficient thermalization and universal quantum computing with quantum gibbs samplers, 2024.

[186] Subir Sachdev. *Quantum Phase Transitions*. Cambridge University Press, 2 edition, 2011.

[187] Subhayan Sahu, Yahui Li, and Pablo Sala. Entanglement cost hierarchies in quantum fragmented mixed states, 2025.

[188] Pablo Sala, Jason Alicea, and Ruben Verresen. Decoherence and wave-function deformation of non-abelian topological order. *Physical Review X*, 15(3), July 2025.

[189] Pablo Sala, Sarang Gopalakrishnan, Masaki Oshikawa, and Yizhi You. Spontaneous strong symmetry breaking in open systems: Purification perspective. *Physical Review B*, 110(15):155150, October 2024.

[190] Pablo Sala and Ruben Verresen. Stability and loop models from decohering non-abelian topological order. *Physical Review Letters*, 134(25), June 2025.

[191] S. Sang, L. A. Lessa, R. S. K. Mong, T. Grover, C. Wang, and T. H. Hsieh. Mixed-state phases from local reversibility. *arXiv preprint*, 2025.

[192] Shengqi Sang and Timothy H. Hsieh. Measurement-protected quantum phases. *Physical Review Research*, 3(2), jun 2021.

[193] Shengqi Sang and Timothy H. Hsieh. Stability of Mixed-State Quantum Phases via Finite Markov Length. *Phys. Rev. Lett.*, 134:070403, Feb 2025.

[194] Shengqi Sang, Leonardo A. Lessa, Roger Mong, Tarun Grover, Chong Wang, and Timothy H. Hsieh, in preparation.

[195] Shengqi Sang, Yaodong Li, Tianci Zhou, Xiao Chen, Timothy H. Hsieh, and Matthew P.A. Fisher. Entanglement negativity at measurement-induced criticality. *PRX Quantum*, 2(3), jul 2021.

[196] Shengqi Sang, Yijian Zou, and Timothy H. Hsieh. Mixed-state quantum phases: Renormalization and quantum error correction. *Phys. Rev. X*, 14:031044, Sep 2024.

[197] K. J. Satzinger, Y.-J Liu, A. Smith, C. Knapp, M. Newman, C. Jones, Z. Chen, C. Quintana, X. Mi, A. Dunsworth, C. Gidney, I. Aleiner, F. Arute, K. Arya, J. Atalaya, R. Babbush, J. C. Bardin, R. Barends, J. Basso, A. Bengtsson, A. Bilmes, M. Broughton, B. B. Buckley, D. A. Buell, B. Burkett, N. Bushnell, B. Chiaro, R. Collins, W. Courtney, S. Demura, A. R. Derk, D. Eppens, C. Erickson, L. Faoro, E. Farhi, A. G. Fowler, B. Foxen, M. Giustina, A. Greene, J. A. Gross, M. P. Harrigan, S. D. Harrington, J. Hilton, S. Hong, T. Huang, W. J. Huggins, L. B. Ioffe, S. V. Isakov, E. Jeffrey, Z. Jiang, D. Kafri, K. Kechedzhi, T. Khattar, S. Kim, P. V. Klimov, A. N. Korotkov, F. Kostritsa, D. Landhuis, P. Laptev, A. Locharla, E. Lucero, O. Martin, J. R. McClean, M. McEwen, K. C. Miao, M. Mohseni, S. Montazeri, W. Mruczkiewicz, J. Mutus, O. Naaman, M. Neeley, C. Neill, M. Y. Niu, T. E. O'Brien, A. Opremcak, B. Pató, A. Petukhov, N. C. Rubin, D. Sank, V. Shvarts, D. Strain, M. Szalay, B. Villalonga, T. C. White, Z. Yao, P. Yeh, J. Yoo, A. Zalcman, H. Neven, S. Boixo, A. Megrant, Y. Chen, J. Kelly, V. Smelyanskiy, A. Kitaev, M. Knap, F. Pollmann, and P. Roushan. Realizing topologically ordered states on a quantum processor. *Science*, 374(6572):1237–1241, 2021.

[198] John Schliemann, J. Ignacio Cirac, Marek Kuś, Maciej Lewenstein, and Daniel Loss. Quantum correlations in two-fermion systems. *Phys. Rev. A*, 64:022303, Jul 2001.

[199] N. Schuch, F. Verstraete, and J. I. Cirac. Nonlocal resources in the presence of superselection rules. *Physical Review Letters*, 92(8), February 2004.

[200] Norbert Schuch and Bela Bauer. Matrix product state algorithms for gaussian fermionic states. *Phys. Rev. B*, 100:245121, Dec 2019.

[201] Norbert Schuch, Frank Verstraete, and J. Ignacio Cirac. Quantum entanglement theory in the presence of superselection rules. *Physical Review A*, 70(4), October 2004.

[202] Benjamin Schumacher and M. A. Nielsen. Quantum data processing and error correction. *Phys. Rev. A*, 54(4):2629–2635, October 1996.

[203] Benjamin Schumacher and Michael D. Westmoreland. Approximate quantum error correction, 2001.

[204] G. Semeghini, H. Levine, A. Keesling, S. Ebadi, T. T. Wang, D. Bluvstein, R. Verresen, H. Pichler, M. Kalinowski, R. Samajdar, A. Omran, S. Sachdev, A. Vishwanath, M. Greiner, V. Vuletić, and M. D. Lukin. Probing topological spin liquids on a programmable quantum simulator. *Science*, 374(6572):1242–1247, 2021.

[205] Hassan Shapourian and Shinsei Ryu. Entanglement negativity of fermions: Monotonicity, separability criterion, and classification of few-mode states. *Physical Review A*, 99(2), February 2019.

[206] Hassan Shapourian and Shinsei Ryu. Finite-temperature entanglement negativity of free fermions. *Journal of Statistical Mechanics: Theory and Experiment*, 2019(4):043106, April 2019.

[207] Hassan Shapourian, Ken Shiozaki, and Shinsei Ryu. Partial time-reversal transformation and entanglement negativity in fermionic systems. *Physical Review B*, 95(16), April 2017.

[208] Nicholas E. Sherman, Trithep Devakul, Matthew B. Hastings, and Rajiv R. P. Singh. Nonzero-temperature entanglement negativity of quantum spin models: Area law, linked cluster expansions, and sudden death. *Physical Review E*, 93(2), February 2016.

[209] Peter W. Shor. Scheme for reducing decoherence in quantum computer memory. *Phys. Rev. A*, 52:R2493–R2496, Oct 1995.

[210] Peter W. Shor. Fault-tolerant quantum computation, 1997.

[211] Brian Skinner, Jonathan Ruhman, and Adam Nahum. Measurement-induced phase transitions in the dynamics of entanglement. *Phys. Rev. X*, 9:031009, Jul 2019.

[212] Ramanjit Sohal and Abhinav Prem. Noisy Approach to Intrinsically Mixed-State Topological Order. *PRX Quantum*, 6:010313, Jan 2025.

[213] C. Spee, K. Schwaiger, G. Giedke, and B. Kraus. Mode entanglement of gaussian fermionic states. *Phys. Rev. A*, 97:042325, Apr 2018.

[214] Adithya Sriram, Tibor Rakovszky, Vedika Khemani, and Matteo Ippoliti. Topology, criticality, and dynamically generated qubits in a stochastic measurement-only kitaev model. *arXiv preprint arXiv:2207.07096*, 2022.

[215] T. M. Stace, S. D. Barrett, and A. C. Doherty. Thresholds for topological codes in the presence of loss. *Phys. Rev. Lett.*, 102:200501, 2009.

[216] E. C. G. Stueckelberg. Quantum Theory in Real Hilbert space. *Helv. Phys. Acta*, 33(8):727–752, 1960.

[217] A. Szabo and N.S. Ostlund. *Modern Quantum Chemistry: Introduction to Advanced Electronic Structure Theory*. Dover Books on Chemistry. Dover Publications, 1996.

[218] B. M. Terhal. Is entanglement monogamous? *IBM Journal of Research and Development*, 48(1):71–78, 2004.

[219] Barbara M. Terhal, Andrew C. Doherty, and David Schwab. Symmetric extensions of quantum states and local hidden variable theories. *Phys. Rev. Lett.*, 90:157903, Apr 2003.

[220] Ruben Verresen, Umberto Borla, Ashvin Vishwanath, Sergej Moroz, and Ryan Thorngren. Higgs condensates are symmetry-protected topological phases: I. discrete symmetries, 2024.

[221] F. Verstraete and J. I. Cirac. Quantum nonlocality in the presence of superselection rules and data hiding protocols. *Physical Review Letters*, 91(1), July 2003.

[222] Frank Verstraete, Michael M. Wolf, and J. Ignacio Cirac. Quantum computation and quantum-state engineering driven by dissipation. *Nature Physics*, 5(9):633–636, 2009.

[223] G. Vidal and R. F. Werner. Computable measure of entanglement. *Physical Review A*, 65(3):032314, February 2002. Publisher: American Physical Society.

[224] Akash Vijay and Jong Yeon Lee. Information critical phases under decoherence, 2025.

[225] Qingyuan Wang, Romain Vasseur, Simon Trebst, Andreas W. W. Ludwig, and Guo-Yi Zhu. Decoherence-induced self-dual criticality in topological states of matter, 2025.

[226] Z. Wang and L. Li. Anomaly in open quantum systems and its implications on mixed-state quantum phases. *arXiv preprint*, 2024.

[227] Zijian Wang, Xu-Dong Dai, He-Ran Wang, and Zhong Wang. Topologically ordered steady states in open quantum systems. *arXiv preprint arXiv:2306.12482*, 2023.

[228] Zijian Wang, Ruihua Fan, Tianle Wang, Samuel J. Garratt, and Ehud Altman. Fractional quantum hall states under density decoherence, 2025.

[229] Zijian Wang and Linhao Li. Anomaly in open quantum systems and its implications on mixed-state quantum phases, March 2024. arXiv:2403.14533 [cond-mat, physics:math-ph, physics:quant-ph].

[230] Zijian Wang, Zhengzhi Wu, and Zhong Wang. Intrinsic mixed-state topological order without quantum memory. *arXiv preprint arXiv:2307.13758*, 2023.

[231] Zijian Wang, Zhengzhi Wu, and Zhong Wang. Intrinsic Mixed-State Topological Order. *PRX Quantum*, 6(1):010314, January 2025.

[232] Xiao-Gang Wen. Topological orders in rigid states. *Int. J. Mod. Phys. B*, 4(2):239–271, 1990.

[233] Xiao-Gang Wen. Quantum Orders in an Exact Soluble Model. *Phys. Rev. Lett.*, 90:016803, Jan 2003.

[234] Reinhard F. Werner. Quantum states with einstein-podolsky-rosen correlations admitting a hidden-variable model. *Phys. Rev. A*, 40:4277–4281, Oct 1989.

[235] Michael M. Wolf, Frank Verstraete, Matthew B. Hastings, and J. Ignacio Cirac. Area laws in quantum systems: Mutual information and correlations. *Physical Review Letters*, 100(7), February 2008.

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

[237] Ting Yu and J. H. Eberly. Sudden death of entanglement. *Science*, 323(5914):598–601, January 2009.

[238] Yìlè Yīng, Maria Ciudad Alañón, Daniel Centeno, Jacopo Surace, Marina Maciel Ansanelli, Ruizhi Liu, David Schmid, and Robert W. Spekkens. On whether quantum theory needs complex numbers: the foil theories perspective, June 2025. arXiv:2506.08091 [quant-ph].

[239] Aidan Zabalo, Michael J. Gullans, Justin H. Wilson, Sarang Gopalakrishnan, David A. Huse, and J. H. Pixley. Critical properties of the measurement-induced transition in random quantum circuits. *Phys. Rev. B*, 101:060301, Feb 2020.

[240] Carolyn Zhang, Yichen Xu, Jian-Hao Zhang, Cenke Xu, Zhen Bi, and Zhu-Xi Luo. Strong-to-weak spontaneous breaking of 1-form symmetry and intrinsically mixed topological order, 2024.

[241] Jian-Hao Zhang, Yang Qi, and Zhen Bi. Strange correlation function for average symmetry-protected topological phases, 2022.

[242] Zhehao Zhang, Utkarsh Agrawal, and Sagar Vijay. Quantum communication and mixed-state order in decohered symmetry-protected topological states, 2024.

[243] Zhehao Zhang, Yijian Zou, Timothy H. Hsieh, and Sagar Vijay. Universal properties of critical mixed states from measurement and feedback, 2025.

[244] Guo-Yi Zhu, Nathanan Tantivasadakarn, Ashvin Vishwanath, Simon Trebst, and Ruben Verresen. Nishimori's cat: stable long-range entanglement from finite-depth unitaries and weak measurements, 2022.

[245] Yijian Zou, Shengqi Sang, and Timothy H. Hsieh. Channeling Quantum Criticality. *Phys. Rev. Lett.*, 130:250403, Jun 2023.

APPENDICES

Appendix A

Supplementary Material for Chapter 2

A.1 Loop patterns and quantum states in physical Hilbert space

Here we establish a correspondence between a loop pattern in the system and the Majorana stabilizer generators within the physical Hilbert space. We emphasize that the loop patterns arise as a consequence of measuring the free Majorana state outside of the physical Hilbert space. However, it is crucial to note that even after these measurements, the resulting states must be projected back into the physical Hilbert space. By utilizing a loop pattern configuration, we demonstrate the construction of stabilizer generators in a canonical form, which allows us to quantify bipartite entanglement by simply counting the number of stabilizer generators connecting the two parts.

A.1.1 Single un-measured boundary

Consider first the scenario of measuring all the bulk qubits of the Toric code, while leaving just a chain of N un-measured spins on the boundary. Using the parton picture, we find that the Majorana fermions on the boundary are paired up using two types of Majorana strands: (A) which connect the nearest neighbor γ_3 and γ_4 Majoranas, and (B) which connect the γ_1 and γ_2 operators through the bulk (see Fig. A.1a).

Note, however, the N type (A) operators don't belong to the physical Hilbert space (set by $D_j = 1$), as they anti-commute with the Projection operators (although, the product of the type (A) operators belong to the stabilizer group as they commute with projection

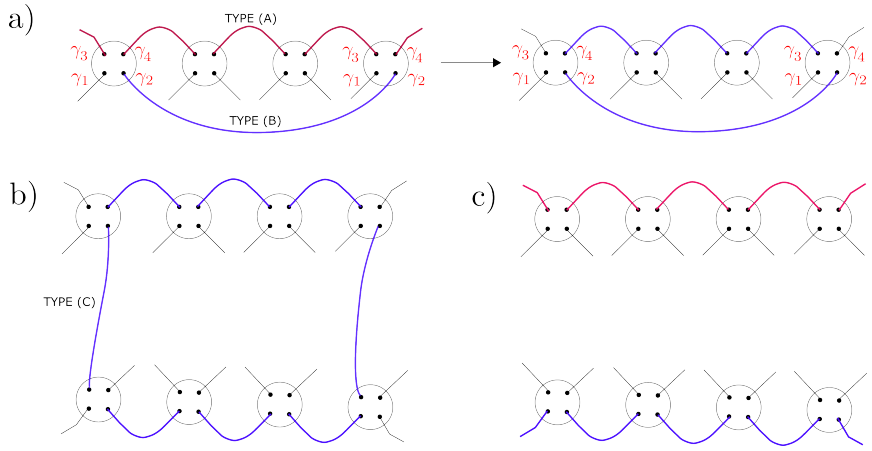


Figure A.1: **One un-measured boundary:**(a) Stabilizer generators of the boundary state from a given pattern of loops can be obtained by multiplying a type (B) pairing (blue strand on the left) with all the interior type (A) strands (colored red in the left). **Two un-measured boundaries:** (b) Stabilizer generator obtained by multiplying two (C) type strands with all intervening (A) strands in upper and lower boundaries. (c) Two independent stabilizer generators which are products of the type (A) strands in upper and lower boundaries.

operators, assuming periodic boundary condition.). However, for each type (B) pairing, we can construct a stabilizer generator for the state by multiplying all the type (A) pairings in its interior, which now commutes with all the on-site projection operators. Since there must be N type (B) pairings, we can form N stabilizer generators, which completely define the state on the boundary. Furthermore, these stabilizer generators are already in the canonical form defined in the main text, i.e. if we take the subset of stabilizer generators with support on any interval and truncate them to that interval, they form an independent set. This implies that if in a particular loop configuration there are n pairings between two parts of the system A and $\overline{A}$, the corresponding quantum state has an entanglement between A and $\overline{A}$ which is $\propto n$.

A.1.2 Two un-measured boundaries

This section focuses on the case where both boundaries of the cylinder are un-measured, and examines the impact of physical Hilbert space projection on the stabilizer generators. Let us call the two boundaries 'upper' and 'lower'. We are interested in calculating the entanglement between the upper and lower boundaries.

As in the previous subsection, there will be the types (A) and (B) pairings on both the boundaries (see Fig A.1b). Again, as before, we can remove the (A) pairings, and form stabilizer generators for each of the (B) type pairings. Note however, none of the (B) type pairings have support on both upper and lower boundaries, and thus don't contribute to any entanglement between them.

In this scenario, there is yet another type of pairing, which pair up $\gamma_{1,2}$ Majoranas between the upper and lower boundaries, which we call type (C). Note, by construction, there can only be even number of such strands. These pairings do not represent physical operators by themselves as they anti-commute with the projection operators. However, we can form physical operators by multiplying two type (C) strands with all the intervening type (A) strands in the upper and lower boundaries. This operator is a physical stabilizer generator, since it intersects all projection operators an even number of times (see Fig. A.1b).

Suppose there are n type (C) pairings in a configuration. If $n = 0$ the entanglement between upper and lower boundaries is clearly zero, as no stabilizer generator in the canonical form has support on both the boundaries. For $n \geq 2$, we can form n such stabilizer generators by the above recipe, as in Fig. A.1b. Note, however, all these generators are not independent. In particular, there are two generators with support only on either upper or lower boundary, given by the product of the the type (A) strands in either boundary (see

Fig. A.1c). This means that we can only form $n-2$ independent stabilizer generators in the canonical form with support on both upper and lower boundaries, leading to $\propto (n-2)$ entanglement between the two boundaries.

A.2 Relation between Majorana partons and Jordan-Wigner fermions

The assignment of Majorana Fermions to spins can be achieved through two different methods, namely the Majorana Parton construction and the Jordan-Wigner transformation. In the former, four Majoranas are assigned to each spin, subject to an additional physical Hilbert space condition ($D_j = \gamma_{j,1}\gamma_{j,2}\gamma_{j,3}\gamma_{j,4} = 1$), where operators that commute with D_j are considered to be in the physical Hilbert space. In contrast, the Jordan-Wigner transformation assigns two Majorana Fermions to each spin. Our objective is to establish a correspondence between certain physical operators in the four Majorana representation and those in the two Majorana representation, when two of the four parton Majoranas are locked in a nearest-neighbor dimer state.

An example of a physical operator in the Majorana parton construction is given by Eq. A.1 (see Fig. A.2a):

$$O_{(j,s),(j',s')} = i\gamma_{j,s} \prod_{k=j'}^{k=j+1} (i\gamma_{k,3}\gamma_{k-1,4})\gamma_{j',s'} \tag{A.1}$$

where s and s' are either 1 or 2. It can be shown that this operator commutes with all D_j's because it has an even number of Majorana fermions labeled by any k. We can express these operators in terms of Pauli operators since they live in a physical Hilbert space. The relation between partons and Pauli operators is given by Eq. A.2.

$$\begin{aligned} X_j &= i\gamma_{j,1}\gamma_{j,2} = i\gamma_{j,4}\gamma_{j,3} \\ Y_j &= i\gamma_{j,2}\gamma_{j,3} = i\gamma_{j,4}\gamma_{j,1} \\ Z_j &= i\gamma_{j,1}\gamma_{j,3} = i\gamma_{j,2}\gamma_{j,4}, \end{aligned} \tag{A.2}$$

Using these relations, we can rewrite the physical operators in Eq. A.1 in terms of Pauli operators, and an example is shown in Eq. A.3.

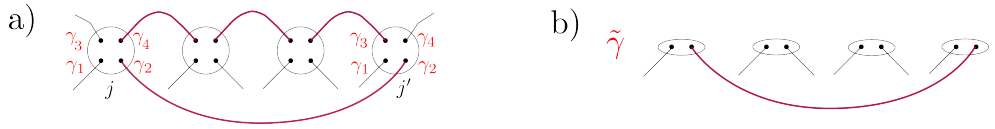


Figure A.2: (a) A physical operator in terms of the parton Majoranas is represented (up to a sign) by the product of all the Majorana strands colored red in the figure. (b) The same operator can be represented as a bilinear operator of Majoranas $\tilde{\gamma}$ obtained by Jordan-Wigner transformation of the qubits. Thus, the Jordan-Wigner transformation of the qubit state is effectively the Majorana state on the $\gamma_{1,2}$ fermions.

$$O_{(j,1),(j',1)} = -Y_j \left(\prod_{k=j+1}^{k=j'-1} X_k \right) Z_{j'} \tag{A.3}$$

Next, we use the Jordan-Wigner transformation to map these Pauli operators to Majorana operators by assigning two Majoranas to each spin. The mapping between Majorana operators and Pauli operators is given by Eq. A.4.

$$\tilde{\gamma}_{j,1} = \left(\prod_{k<j} X_k \right) Z_j, \;\; \tilde{\gamma}_{j,2} = \left(\prod_{k<j} X_k \right) Y_j \tag{A.4}$$

Using this transformation, we can rewrite the operators in Eq. A.3 in terms of Majorana operators, as shown in Eq. A.5.

$$Y_j \left(\prod_{k=j+1}^{k=j'-1} X_k \right) Z_{j'} = -i\tilde{\gamma}_{j,1}\tilde{\gamma}_{j',1} \tag{A.5}$$

Using this mapping, any general physical operators in Eq. A.1 can be represented as bilinears of $\tilde{\gamma}$ Majoranas,

$$i\gamma_{j,s} \prod_{k=j'}^{k=j+1} (i\gamma_{k,3}\gamma_{k-1,4})\gamma_{j',s'} = i\tilde{\gamma}_{j,s}\tilde{\gamma}_{j',s'} \tag{A.6}$$

Here, s and s' are either 1 or 2. This shows that by removing γ_3 and γ_4 from any physical operator, we obtain the Jordan-Wigner transformed version of that operator, as shown in Fig. A.2b. Thus, the Jordan-Wigner transformation of the qubit state is effectively the Majorana state on the $\gamma_{1,2}$ fermions. Going the other direction, we can read out the qubit operator by performing a standard Jordan-Wigner transformation on the operator involving the $\gamma_{1,2}$ Majoranas only.

A.3 Jordan-Wigner transformation of the boundary after bulk measurements

In this section, we generalize the results from Appendix A.2 from stabilizer states to general Gaussian states of Majorana operators.

We start with the density matrix of the final Gaussian state before projections and determine how the projections affect it. Any general Gaussian state can be represented by an orthogonal transformation of the original Majorana fermions γ, conveniently arranged as a vector $\vec{\gamma} = (\gamma_{1,1}, \gamma_{1,2}, \gamma_{2,1}, ..., \gamma_{2,n})^T$. Let us define $\vec{\gamma}'$ using a generic orthogonal matrix O as $\vec{\gamma}' = O\vec{\gamma}$.

It is straightforward to show that $\tilde{\gamma}$ also satisfies the Majorana algebra. The most general boundary density matrix in our setup is composed of two parts: the first part comes from the initial short stabilizers on the boundary, and the second part comes from the Gaussian operations in the bulk. Such a general density matrix can be represented as [26],

$$\rho = \prod_{j=1}^{j=N-1} \frac{1 + i\gamma_{j+1,3}\gamma_{j,4}}{2} \prod_{k=1}^{k=N} \frac{1 + s_k(i\gamma'_{2k}\gamma'_{2k+1})}{2}, \tag{A.7}$$

where $s_k \in [-1, 1]$.

We need to apply the projections projections $P_j = (1 + \gamma_{j,1}\gamma_{j,2}\gamma_{j,3}\gamma_{j,4})/2$ to the density matrix to get a physical density matrix,

$$\rho_f = \left(\prod_{j=1}^{j=N} P_j\right) \rho \left(\prod_{j=1}^{j=N} P_j\right).$$

We can re-express $i\gamma'_{2k}\gamma'_{2k+1}$ in terms of the original Majorana fermions as $\sum_{i,j} a^k_{i,j} i\gamma_i\gamma_j$, where $a^k_{i,j} = O_{2k,i}O_{2k+1,j}$. It is now straightforward to show that Eq. A.7 can be rewritten

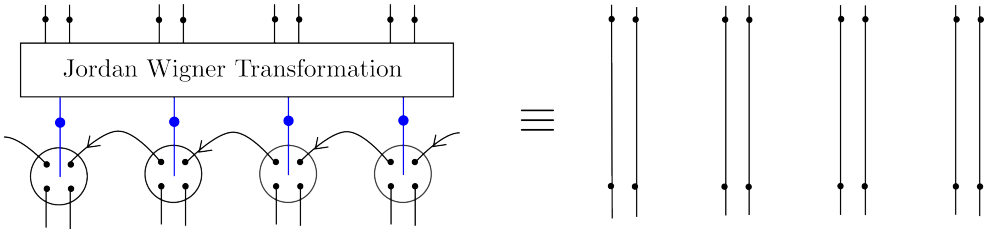


Figure A.3: The key result of this section is the cancellation between the projection tensor applied on the dimer state and the Jordan-Wigner transformation. Each large circle represents a projection tensor as defined in Chapter 2. Two of the four Majoranas per circle are paired in the depicted nearest neighbor dimer state. By contracting these tensors in this manner, we obtain a tensor that maps $2N$ Majorana fermions (bottom) to N spins (blue dots). Subsequently, the Jordan-Wigner tensor takes N spins as input and outputs $2N$ Majorana fermions. These two tensors cancel each other out when combined.

as,

$$\rho = \prod_{j=1}^{j=N} \frac{1 + i\gamma_{j,3}\gamma_{j+1,4}}{2} \prod_{k=1}^{k=N} \frac{1 + s_k \sum_{i,j} a_{i,j}^k i\gamma_i (\prod_{k=j}^{k=i+1} (i\gamma_{k,3}\gamma_{k-1,4}))\gamma_j}{2}, \tag{A.8}$$

where we have introduced a Majorana string in the second factor, which can always be absorbed in the first factor.

We introduce a new notation to represent the density matrix as the product of two parts $\rho = \rho_1 \rho_2$,

$$\rho_1 = \prod_{j=1}^{j=N} \frac{1+i\gamma_{j,3}\gamma_{j+1,4}}{2}$$
$$\rho_2 = \prod_{k=1}^{k=N} \frac{1+s_k \sum_{i,j} a_{i,j}^k i\gamma_i (\prod_{k=j}^{k=i+1} (i\gamma_{k,3}\gamma_{k-1,4}))\gamma_j}{2}.$$

One can easily check that ρ_2 and P_j commute, due to the fact that ρ_2 intersects even number of times in each factor of P_j. This allows us to represent the density matrix after projection as $\rho_f = (\prod P_j)\, \rho_1\, (\prod P_j)\, \rho_2$. The computation of $(\prod P_j)\, \rho_1\, (\prod P_j)$ is now amenable to stabilizer formalism. We can directly borrow the results of Appendix A.2 to re-express the parton Majoranas in terms of the Jordan Wigner Majoranas, $\gamma \to \tilde{\gamma}$. As described in that Appendix, we can simply remove the γ_3 and γ_4 operators to represent a physical density matrix after projection in terms of the Jordan Wigner Majoranas $\tilde{\gamma}$,

$$\rho_f = \prod_{k=1}^{k=N} \frac{1 + s_k \sum_{i,j} a_{i,j}^k i\tilde{\gamma}_i \tilde{\gamma}_j}{2}. \tag{A.9}$$

The $\tilde{\gamma}$ Majoranas can now be rotated using the orthogonal transformation O to get back to the initial Gaussian state we began with in terms of $\tilde{\gamma}'$ Majoranas,

$$\rho_f = \prod_{k=1}^{k=N} \frac{1 + s_k(i\tilde{\gamma}'_{2k}\tilde{\gamma}'_{2k+1})}{2}. \tag{A.10}$$

These series of manipulations can be depicted simply in Fig. A.3. In summary, the boundary projections followed by a Jordan Wigner transformation lead us back to the original Gaussian state before enforcing the projections.

Appendix B

Supplementary Material for Chapter 3

B.1 Feedback Channel

Here, we describe the feedback channel F used in the main text. The feedback channel is an important step in performing MBQC, since the measurement outcomes of the code qubits are random. If the measurements of the code qubits are postselected to be in the $|+\rangle$ state, then the state of the code qubits is seamlessly teleported to the final layer. Otherwise, if one or more code qubits are measured in the $|-\rangle$ state, the realized circuit deviates from the intended one by byproduct operators [see Eq. (3.2)], which must be tracked and corrected via the feedback channel F. We argue here that the feedback channel allows us to map the measured resource state to the circuit model and a map F^{-1} takes us back from the circuit model to the measured resource state. We start by considering the noiseless case.

B.1.1 Feedback Channel for Noiseless MBQC

We begin by introducing F in the absence of any noise. Using the notation in the main text, we define the density matrix of the circuit in the noiseless limit as $\tilde{\rho}_0$. Given an initial code state ρ_i, measuring a syndrome and recording its outcome is described by the channel M, acting as follows:

$$M(\rho_i) = \sum_s P_s \rho_i P_s \otimes |s\rangle\langle s|, \tag{B.1}$$

where P_s is the projection operator associated with syndrome measurement outcome s, and $|s\rangle\langle s|$ denotes the syndrome state. Extending this process over multiple instances of

syndrome extraction, the final density matrix $\tilde{\rho}_0$ is:

$$\tilde{\rho}_0 = \sum_{\vec{S}} P_{\vec{S}} \rho_i P_{\vec{S}} \otimes |\vec{S}\rangle\langle\vec{S}|, \tag{B.2}$$

where $\vec{S}$ represents the sequence of all syndrome outcomes, and $P_{\vec{S}}$ is the corresponding projection operator sequence i.e., the product of projectors $P_{\vec{S}} = P_{\vec{S}_{t_f}} \cdots P_{\vec{S}_1}$. The state $\tilde{\rho}_0$ both encodes the final state of the circuit as well as all syndrome qubits through the history. Note that in the absence of errors, for a stabilizer code, as opposed to a subsystem code, all of the syndrome outcomes are 0.

In contrast, the post-measurement density matrix of the MBQC resource state can be expressed as:

$$\rho_0 = \frac{1}{2^{|\vec{C}|}} \sum_{\vec{C}} |\vec{C}\rangle\langle\vec{C}| \otimes \tilde{\rho}_{0,\vec{C}}, \tag{B.3}$$

where $|\vec{C}\rangle\langle\vec{C}|$ corresponds to the outcomes of measuring code qubits, each equally likely. The conditional state $\tilde{\rho}_{0,\vec{C}}$ encodes both syndrome measurement outcomes and the final boundary state of the code qubits. Note that, in this case, the syndrome outcomes and state on the final layer depend on the measurement outcomes of the code qubits because of the Pauli byproducts associated with $|-\rangle$ measurements of the code qubits. Thus, $\tilde{\rho}_{0,\vec{C}}$ can be written as:

$$\tilde{\rho}_{0,\vec{C}} = \sum_{\vec{S}} P_{\vec{S}_{t_f}} \Big(\dots U_{\vec{C}_3} \Big(P_{\vec{S}_2} \big(U_{\vec{C}_2} \left(P_{\vec{S}_1} \left(U_{\vec{C}_1} \rho_i U_{\vec{C}_1}^\dagger \right) P_{\vec{S}_1} \right) U_{\vec{C}_2}^\dagger \big) P_{\vec{S}_2} \Big) U_{\vec{C}_3}^\dagger \dots \Big) P_{\vec{S}_{t_f}} \otimes |\vec{S}\rangle\langle\vec{S}|, \tag{B.4}$$

where $U_{\vec{C}_i}$ denotes the Pauli byproducts at a given time step i determined by the code qubit measurements, as written in Eq. (3.2), and $P_{\vec{S}_j}$ are projectors associated with syndrome measurements at each time step. From this expression, we now construct the feedback channel, which involves commuting the Pauli byproducts from the innermost parentheses to the outermost.

Since the $U_{\vec{C}_i}$ operators are Pauli operators, commuting them through the projectors $P_{\vec{S}_j}$ in Eq. (B.4) can, at most, flip the signs of the components of $\vec{S}_j$

$$\tilde{\rho}_{0,\vec{C}} = \sum_{\vec{S}} \Big(\prod_{\vec{C}_i} U_{\vec{C}_i} \Big) P_{\vec{S}'_{t_f}} \left(\dots \left(P_{\vec{S}'_2} \left(P_{\vec{S}'_1} \rho_i P_{\vec{S}'_1} \right) P_{\vec{S}'_2} \right) \dots \right) P_{\vec{S}'_{t_f}} \Big(\prod_{\vec{C}_i} U_{\vec{C}_i} \Big)^\dagger \otimes |\vec{S}\rangle\langle\vec{S}| \tag{B.5}$$

By changing the summing variable from $\vec{S}$ to $\vec{S}'$ we get the following

$$\tilde{\rho}_{\vec{C},p} = \sum_{\vec{S}} (\prod_{\vec{C}_i} U_{\vec{C}_i}) P_{\vec{S}_{t_f}} (\dots (P_{\vec{S}_2} (P_{\vec{S}_1} \rho_i P_{\vec{S}_1}) P_{\vec{S}_2}) \dots) P_{\vec{S}_{t_f}} (\prod_{\vec{C}_i} U_{\vec{C}_i})^\dagger \otimes |\vec{S}'\rangle\langle\vec{S}'| \tag{B.6}$$

which effectively absorbs the sign changes into the $|\vec{S}\rangle\langle\vec{S}|$ term. The state $|\vec{S}'\rangle\langle\vec{S}'|$ can be transformed back to state $|\vec{S}\rangle\langle\vec{S}|$ by a unitary transformation which results in the final form of:

$$\tilde{\rho}_{0,\vec{C}} = U_{\vec{C}} \left(\sum_{\vec{S}} P_{\vec{S}_{t_f}} (\dots (P_{\vec{S}_2} (P_{\vec{S}_1} \rho_i P_{\vec{S}_1}) P_{\vec{S}_2}) \dots) P_{\vec{S}_{t_f}} \otimes |\vec{S}\rangle\langle\vec{S}| \right) U_{\vec{C}}^\dagger. \tag{B.7}$$

where the unitary acts on the code qubits as the product of the $U_{\vec{C}_i}$'s and on the syndromes.

Consequently, the full resource state is:

$$\rho_0 = \frac{1}{2^{|\vec{C}|}} \sum_{\vec{C}} |\vec{C}\rangle\langle\vec{C}| \otimes U_{\vec{C}} (\tilde{\rho}_0) U_{\vec{C}}^\dagger, \tag{B.8}$$

where $\tilde{\rho}_0$ is the postselected MBQC state equivalent to the circuit state, including syndromes and outputs.

This establishes a quantum channel F relating the MBQC state ρ_0 to its equivalent circuit representation $\tilde{\rho}_0$:

$$F(\rho_0) = \tilde{\rho}_0, \quad F^{-1}(\tilde{\rho}_0) = \rho_0, \tag{B.9}$$

explicitly given by:

$$F(\cdot) = \mathrm{Tr}_C \left(\sum_{\vec{C}} U_{\vec{C}}^\dagger P_{\vec{C}} (\cdot) P_{\vec{C}} U_{\vec{C}} \right), \quad F^{-1}(\cdot) = \frac{1}{2^{|\vec{C}|}} \sum_{\vec{C}} |\vec{C}\rangle\langle\vec{C}| \otimes \left(U_{\vec{C}} (\cdot) U_{\vec{C}}^\dagger \right), \tag{B.10}$$

where $P_{\vec{C}}$ projects onto code qubit measurement outcomes and $U_{\vec{C}}$ are the associated non-adaptive Pauli corrections. Note that F^{-1} serves as a right inverse, satisfying $F \circ F^{-1} = \mathbb{I}$.

B.1.2 Feedback Channel for Noisy MBQC

We now generalize the relationship in Eq. (B.10) to noisy settings. That is, if $\tilde{\rho}_p$ denotes the circuit state with noise parameter p and ρ_p the corresponding noisy MBQC state (defined via the noise map in the main text [Fig. 3.3]), the feedback channel still holds:

$$F(\rho_p) = \tilde{\rho}_p, \quad F^{-1}(\tilde{\rho}_p) = \rho_p. \tag{B.11}$$

For a noisy circuit, the density matrix is constructed by sequentially applying the measurement channel and the noise channel, expressed as:

$$\tilde{\rho}_p = \left(\circ_{t=1}^{t_f} (M \circ \mathcal{E}_p) \right) (\rho_i), \tag{B.12}$$

where $\mathcal{E}_p$ denotes the noise channel acting on both the code state and syndrome outcomes, incorporating errors such as X- and Z-type noise, as well as readout errors. Explicitly, the noisy state $\tilde{\rho}_p$ is

$$\tilde{\rho}_p = \sum_{\vec{S}} P_{\vec{S}_{t_f}} \left(\cdots \mathcal{E}_p \left(P_{\vec{S}_2} \left(\mathcal{E}_p \left(P_{\vec{S}_1} \left(\mathcal{E}_p(\rho_i) \right) P_{\vec{S}_1} \right) \right) P_{\vec{S}_2} \right) \cdots \right) P_{\vec{S}_{t_f}} \otimes \mathcal{E}_p(|\vec{S}\rangle\langle\vec{S}|). \tag{B.13}$$

To prove Eq. (B.11), it suffices to show:

$$F^{-1}(\tilde{\rho}_p) = \frac{1}{2^{|\vec{C}|}} \sum_{\vec{C}} |\vec{C}\rangle\langle\vec{C}| \otimes (U_{\vec{C}} \tilde{\rho}_p U_{\vec{C}}^\dagger) = \rho_p, \tag{B.14}$$

The identity $F \circ F^{-1} = \mathbb{I}$ then implies the converse: $F(\rho_p) = \tilde{\rho}_p$.

First, consider the noiseless case ($p = 0$):

$$\rho_0 = \frac{1}{2^{|\vec{C}|}} \sum_{\vec{C}} |\vec{C}\rangle\langle\vec{C}| \otimes \tilde{\rho}_{0,\vec{C}}, \tag{B.15}$$

Applying the $\mathcal{E}_p^Z$ channel yields:

$$\rho_p = \frac{1}{2^{|\vec{C}|}} \sum_{\vec{C}} \mathcal{E}_p^Z(|\vec{C}\rangle\langle\vec{C}| \otimes \tilde{\rho}_{0,\vec{C}}), \tag{B.16}$$

Let us start by applying the noise channel on only one code qubit (say C_1), we then have:

$$\begin{aligned}
\mathcal{E}_p^{C_1}(\rho_0) &= \frac{1}{2^{|\vec{C}|}} \sum_{\vec{C}} \left(|\vec{C}_1'\rangle\langle\vec{C}_1'| \right) \otimes \mathcal{E}_p^z \left(|C_1\rangle\langle C_1| \right) \otimes \tilde{\rho}_{0,\vec{C}} \qquad \text{(B.17)} \\
&= \frac{1}{2^{|\vec{C}|}} \sum_{\vec{C}} \left(|\vec{C}_1'\rangle\langle\vec{C}_1'| \right) \otimes \left((1-p)|C_1\rangle\langle C_1| + p|-C_1\rangle\langle -C_1| \right) \otimes \tilde{\rho}_{0,\vec{C}} \\
&= \frac{1}{2^{|\vec{C}|}} \sum_{\vec{C}} |\vec{C}\rangle\langle\vec{C}| \otimes \left((1-p)\tilde{\rho}_{0,\{\vec{C}_1', C_1\}} + p\, \tilde{\rho}_{0,\{\vec{C}_1', -C_1\}} \right),
\end{aligned}$$

Here, the vector $\vec{C}$ is broken into two parts: $\vec{C}_1'$ and C_1, where $\vec{C}_1'$ denotes all code qubits except C_1. By relabeling the code basis states, the net effect is to move the noise action from the code qubits $|\vec{C}\rangle\langle\vec{C}|$ onto $\tilde{\rho}_{0,\vec{C}}$.

Now, we write the noise on $\tilde{\rho}_{0,\vec{C}}$ explicitly:

$$\begin{aligned}&\Big((1-p)\tilde{\rho}_{0,\{\vec{C}_1',C_1\}} + p\ \tilde{\rho}_{0,\{\vec{C}_1',-C_1\}}\Big) \\ &= \sum_{\vec{S}} P_{\vec{S}_{t_f}} \Big(... U_{\vec{C}_3}\Big(P_{\vec{S}_2}\Big(U_{\vec{C}_2}\Big(P_{\vec{S}_1}U_{\vec{C}_1}\mathcal{E}_p(\rho_i)U_{\vec{C}_1}^\dagger P_{\vec{S}_1}\Big)U_{\vec{C}_2}^\dagger\Big)P_{\vec{S}_2}\Big)U_{\vec{C}_3}^\dagger ...\Big)P_{\vec{S}_{t_f}} \otimes (|\vec{S}\rangle\langle\vec{S}|),\end{aligned} \tag{B.18}$$

where the noise channel $\mathcal{E}_p$ is $\mathcal{E}_p(\cdot) = (1-p)(\cdot) + p\ U_{C_1}U_{-C_1}(\cdot)U_{-C_1}U_{C_1}$, with the byproduct operators U_{C_1}, U_{-C_1} in Eq. (3.2). Note that the factor of U_{C_1} in $U_{\vec{C}_1}$ cancels with the factor of U_{C_1} in the action of $\mathcal{E}_p$. Extending this procedure to all qubits yields the explicit form:

$$\begin{aligned}\rho_p &= \frac{1}{2^{|\vec{C}|}}\sum_{\vec{C}} |\vec{C}\rangle\langle\vec{C}| \otimes \tilde{\rho}_{\vec{C},p}, \\ \tilde{\rho}_{\vec{C},p} &= \sum_{\vec{S}} P_{\vec{S}_{t_f}}\Big(... U_{\vec{C}_3}\mathcal{E}_p\Big(P_{\vec{S}_2}\Big(U_{\vec{C}_2}\mathcal{E}_p\Big(\\ &\qquad P_{\vec{S}_1}(U_{\vec{C}_1}\mathcal{E}_p(\rho_i)U_{\vec{C}_1}^\dagger)P_{\vec{S}_1}\Big)U_{\vec{C}_2}^\dagger\Big)P_{\vec{S}_2}\Big)U_{\vec{C}_3}^\dagger ...\Big)P_{\vec{S}_{t_f}} \otimes \mathcal{E}_p(|\vec{S}\rangle\langle\vec{S}|).\end{aligned} \tag{B.19}$$

Now we aim to commute the $U_{\vec{C}_i}$ operators through the $\mathcal{E}_p$'s and the $P_{\vec{S}_i}$'s, moving it from inside the parentheses to the outermost position. First, note that $U_{\vec{C}_i}$ commutes with any $\mathcal{E}_p$, or equivalently: $\mathcal{E}_p(U_{\vec{C}_i}(\cdot)U_{\vec{C}_i}^\dagger) = U_{\vec{C}_i}\mathcal{E}_p(\cdot)U_{\vec{C}_i}^\dagger$. Second, $U_{\vec{C}_j}$ either commutes with $P_{\vec{S}_i}$ or changes the sign of S_i. We define $\vec{S}_i'$ according to: $P_{\vec{S}_i}(U_{\vec{C}_j}(\cdot)U_{\vec{C}_j}^\dagger)P_{\vec{S}_i}^\dagger = U_{\vec{C}_j}(P_{\vec{S}_i'}(\cdot)P_{\vec{S}_i'}^\dagger)U_{\vec{C}_j}^\dagger$. Using these two commutation relations, we obtain:

$$\tilde{\rho}_{\vec{C},p} = \sum_{\vec{S}}(\prod_{C_i} U_{C_i})P_{\vec{S}_{t_f}'}\Big(... \mathcal{E}_p\Big(P_{\vec{S}_2'}\Big(\mathcal{E}_p\Big(P_{\vec{S}_1'}\mathcal{E}_p(\rho_i)P_{\vec{S}_1'}\Big)\Big)P_{\vec{S}_2'}\Big)...\Big)P_{\vec{S}_{t_f}'}(\prod_{C_i} U_{C_i})^\dagger \otimes \mathcal{E}_p(|\vec{S}\rangle\langle\vec{S}|) \tag{B.20}$$

Given that $|\vec{S}\rangle\langle\vec{S}|$ can be expressed as a correction applied to $|\vec{S}'\rangle\langle\vec{S}'|$ with a unitary feedback that flips the relevant syndromes, we obtain:

$$\tilde{\rho}_{\vec{C},p} = U_{\vec{C}}\left(\sum_{\vec{S}'} P_{\vec{S}_{t_f}'}\Big(... \mathcal{E}_p\Big(P_{\vec{S}_2'}\Big(\mathcal{E}_p\Big(P_{\vec{S}_1'}\mathcal{E}_p(\rho_i)P_{\vec{S}_1'}\Big)\Big)P_{\vec{S}_2'}\Big)...\Big)P_{\vec{S}_{t_f}'} \otimes \mathcal{E}_p(|\vec{S}'\rangle\langle\vec{S}'|)\right)U_{\vec{C}}^\dagger. \tag{B.21}$$

Inserting Eq. (B.21) into Eq. (B.13), we obtain:

$$\tilde{\rho}_{\vec{C},p} = U_{\vec{C}} \tilde{\rho}_p U_{\vec{C}}^\dagger. \tag{B.22}$$

Finally, we arrive at:

$$\rho_p = \frac{1}{2^{|\vec{C}|}} \sum_{\vec{C}} |\vec{C}\rangle\langle\vec{C}| \otimes (U_{\vec{C}} \tilde{\rho}_p U_{\vec{C}}^\dagger), \tag{B.23}$$

which completes the proof.

B.2 Coherent errors in the syndrome extraction circuit

In the main text, we established a connection between the threshold of the syndrome extraction circuit and a mixed-state phase transition under incoherent noise. In this appendix, we consider instead coherent errors in the syndrome extraction circuit and use two complementary approaches. We first consider the F channel from Appendix B.1 and discuss its subtleties and limitations in the presence of coherent errors. We nevertheless demonstrate that for the simple case of the repetition code with coherent errors acting only on the code qubits, this formalism successfully establishes a connection between mixed-state phase transitions and the fault-tolerance threshold.

Next, we use measurement and postselection on the resource state to show that the CMI of the syndromes in the circuit is equal to that in the MBQC model in the presence of coherent errors. This result is particularly useful, as it enables a correspondence whereby the divergence of the Markov length in one setup implies the divergence of the Markov length in the other. However, it is not yet strong enough to establish that the threshold under coherent noise corresponds to a mixed-state phase transition.

B.2.1 Coherent error and F channel

Here we study whether the fault-tolerance transition observed in a syndrome extraction circuit subjected to coherent errors can be mapped to a mixed-state phase transition of the X-dephased resource state. To establish this correspondence using the arguments in the main text, we must demonstrate that the mixed state of the circuit under coherent errors, denoted by $\tilde{\rho}_\theta$, maps via the inverse transformation $F^{-1}(\tilde{\rho}_\theta)$ to a local Lindbladian evolution acting upon a completely decohered resource state ρ_0:

$$F^{-1}(\tilde{\rho}_\theta) = \mathcal{T} e^{\int_0^{t_p} \mathcal{L}_1(t) dt}[\rho_0]. \tag{B.24}$$

For simplicity, we consider the case where the coherent noise acts only on the code qubits.

We first demonstrate explicitly how $\tilde{\rho}_\theta$ decomposes into a combination of correlated incoherent errors and a residual term that we expect to vanish in the thermodynamic limit. We show vanishing of the residual term explicitly for the repetition code syndrome extraction circuit. Subsequently, using the transformation F^{-1}, correlated incoherent errors in the circuits maps to correlated incoherent errors within the resource state. We further verify explicitly in the repetition-code case that such errors correspond precisely to a local Lindbladian evolution. However, the generality of this construction beyond the repetition-code setting remains an open question.

The evolution of the density matrix is modeled as a sequence of alternating noise and measurement channels:

$$\rho_f = \left(\circ_{t=1}^{t_f}(M \circ \mathcal{E})\right)(\rho_i), \tag{B.25}$$

where $\mathcal{E}$ denotes the noise channel that acts with coherent noise on the code qubits and readout errors on the syndrome outcomes. The final state of the system, after extracting syndromes $\vec{S} = (\vec{S}_1, \vec{S}_2, \ldots, \vec{S}_{t_f})$, is given by:

$$\rho_f = \sum_{\vec{S}} P_{\vec{S}_{t_f}} \left(\cdots \mathcal{E}\left(P_{\vec{S}_2}\left(\mathcal{E}\left(P_{\vec{S}_1}\left(\mathcal{E}(\rho_i)\right) P_{\vec{S}_1}\right)\right) P_{\vec{S}_2}\right) \cdots\right) P_{\vec{S}_{t_f}} \otimes \mathcal{E}_p(|\vec{S}\rangle\langle\vec{S}|), \tag{B.26}$$

where $P_{\vec{S}_t}$ are the syndrome projection operators and $\mathcal{E}_p$ is the readout error channel on the classical syndrome record.

From here onward, we assume the noise channel $\mathcal{E}$ acts as a coherent unitary rotation $e^{i\alpha\hat{X}} = \cos(\alpha) + i\sin(\alpha)\hat{X}$ on code qubits and the underling code is a stabilizer code. The action of this channel on the initial state ρ_i can be expanded as:

$$\mathcal{E}(\rho_i) = \sum_{\vec{m}} C_{\vec{m}} \hat{X}^{\vec{m}_L} \rho_i \hat{X}^{\vec{m}_R}, \tag{B.27}$$

where $\vec{m}_L$ and $\vec{m}_R$ denote binary vectors indexing Pauli-X operators on each qubit, therfore $\hat{X}^{\vec{m}} = \hat{X}^{\vec{m}_1}\hat{X}^{\vec{m}_2}\cdots\hat{X}^{\vec{m}_n}$. $C_{\vec{m}}$ are coefficients determined by products of $\cos(\alpha)$ and $i\sin(\alpha)$. Now we move to the next parenthesis of Eq. (B.26) which projects to the syndrome outcome $\vec{S}_1$:

$$P_{\vec{S}_1}\mathcal{E}(\rho_i)P_{\vec{S}_1} = \sum_{\vec{m}} C_{\vec{m}} P_{\vec{S}_1} \hat{X}^{\vec{m}_L} \rho_i \hat{X}^{\vec{m}_R} P_{\vec{S}_1}. \tag{B.28}$$

This can be simplified taking the initial state to be $\rho_i = \rho_L P_{\vec{S}=0}$, where $P_{\vec{S}=0}$ projects in to the code subspace and ρ_L further chooses a specific code state:

$$P_{\vec{S}_1}\hat{X}^{\vec{m}_L}\rho_i\hat{X}^{\vec{m}_R}P_{\vec{S}_1} = P_{\vec{S}_1}\hat{X}^{\vec{m}_L}\rho_L P_{\vec{S}=0}\hat{X}^{\vec{m}_R}P_{\vec{S}_1} = \rho_L' P_{\vec{S}_1}P_{\vec{S}_0}\hat{X}^{\vec{m}_L}\hat{X}^{\vec{m}_R}P_{\vec{S}_1}, \tag{B.29}$$

where $P_{\vec{S}_0} = \hat{X}^{\vec{m}_L} P_{\vec{S}=0} \hat{X}^{\vec{m}_L}$ and $\rho'_L = \hat{X}^{\vec{m}_L} \rho_L \hat{X}^{\vec{m}_L}$. Now note that if $\hat{X}^{\vec{m}_L} \hat{X}^{\vec{m}_R}$ does not commute with $P_{\vec{S}_1}$, the $P_{\vec{S}_1} P_{\vec{S}_0} \hat{X}^{\vec{m}_L} \hat{X}^{\vec{m}_R} P_{\vec{S}_1}$ will be zero. This is because by defining $P_{\vec{S}'_1} = \hat{X}^{\vec{m}_L} \hat{X}^{\vec{m}_R} P_{\vec{S}_1} \hat{X}^{\vec{m}_L} \hat{X}^{\vec{m}_R}$ the following equation holds since the projections are orthogonal.

$$P_{\vec{S}_1} P_{\vec{S}_0} P_{\vec{S}'_1} = 0, \tag{B.30}$$

Thus, only terms for which $\hat{X}^{\vec{m}_L} \hat{X}^{\vec{m}_R}$ commutes with $P_{\vec{S}_1}$ survive. These terms correspond to logical operators or stabilizers. We can decompose the sum in Eq. (B.28) into three contributions:

$$\begin{aligned} P_{\vec{S}_1} \mathcal{E}(\rho_i) P^\dagger_{\vec{S}_1} = & \sum_{\vec{m}_L} C^1_{\vec{m}_L} P_{\vec{S}_1} \hat{X}^{\vec{m}_L} \rho_i \hat{X}^{\vec{m}_L} P_{\vec{S}_1} \\ & + \sum_S \sum_{\vec{m}_L} C^2_{\vec{m}_L} P_{\vec{S}_1} \hat{X}^{\vec{m}_L} \rho_i X^S \hat{X}^{\vec{m}_L} P_{\vec{S}_1} \\ & + \sum_{\vec{m}_L} D_{\vec{m}_L} P_{\vec{S}_1} \hat{X}^{\vec{m}_L} \rho_i \hat{X}^{\vec{m}_L} \bar{X} P_{\vec{S}_1}, \end{aligned} \tag{B.31}$$

where X^S are stabilizers and $\bar{X}$ denotes a logical Pauli-X operator. Since ρ_i is stabilized by X^S, the second term can be absorbed into the first. We are left with:

$$\begin{aligned} P_{\vec{S}_1} \mathcal{E}(\rho_i) P^\dagger_{\vec{S}_1} = & \sum_{\vec{m}_L} C_{\vec{m}_L} P_{\vec{S}_1} \hat{X}^{\vec{m}_L} \rho_i \hat{X}^{\vec{m}_L} P_{\vec{S}_1} \\ & + \sum_{\vec{m}_L} D_{\vec{m}_L} P_{\vec{S}_1} \hat{X}^{\vec{m}_L} \rho_i \hat{X}^{\vec{m}_L} \bar{X} P_{\vec{S}_1}. \end{aligned} \tag{B.32}$$

We define maps as the following:

$$\mathcal{N}_{\text{sym}}(\rho_i) = \sum_{\vec{m}_L} C_{\vec{m}_L} \hat{X}^{\vec{m}_L} \rho_i \hat{X}^{\vec{m}_L}, \tag{B.33}$$

$$\mathcal{N}_{\text{asym}}(\rho_i) = \sum_{\vec{m}_L} D_{\vec{m}_L} \hat{X}^{\vec{m}_L} \rho_i \hat{X}^{\vec{m}_L} \bar{X}. \tag{B.34}$$

Thus, the total channel can be expressed as:

$$P_{\vec{S}_1} \mathcal{E}(\rho_i) P_{\vec{S}_1} = P_{\vec{S}_1} \left(\mathcal{N}_{\text{sym}}(\rho_i) + \mathcal{N}_{\text{asym}}(\rho_i) \right) P_{\vec{S}_1}. \tag{B.35}$$

The symmetric part, $\mathcal{N}_{\text{sym}}$, corresponds to correlated incoherent noise and can be mapped to a corresponding error model in MBQC using the construction introduced in

the main text. The asymmetric part, $\mathcal{N}_{\text{asym}}$, does not follow this mapping. However, if one is interested only in mixed-state properties—which should be independent of the logical encoding—one can either select a suitable ρ_i to further symmetrize the asymmetric term or show that $\mathcal{N}_{\text{asym}}$ is suppressed in the thermodynamic limit. In the next subsection, we explicitly demonstrate this suppression for the case of the repetition code. And we further show the Eq. (B.24) holds for the repetition code by showing that $\mathcal{N}_{\text{asym}}$ is a local channel, however the locality of this channel for the general setting remains an open question.

B.2.2 Coherent error in repetition code

In the case of the repetition code, $\hat{X}^{\vec{m}_L}$ and $\hat{X}^{\vec{m}_R}$ must satisfy $\hat{X}^{\vec{m}_L}\hat{X}^{\vec{m}_R} = I$ or $\bar{X}$, since there is no X-type stabilizer. This simplifies the following sum:

$$P_{\vec{S}_1}\left(\mathcal{E}(\rho_i)\right)P_{\vec{S}_1}^{\dagger} = \sum_{\vec{m}_L} C^1_{\vec{m}_L} P_{\vec{S}_1}\left(\hat{X}^{\vec{m}_L}(\rho_i)\hat{X}^{\vec{m}_L}\right)P_{\vec{S}_1} + \sum_{\vec{m}_L} C^2_{\vec{m}_L} P_{\vec{S}_1}\left(\hat{X}^{\vec{m}_L}(\rho_i)\hat{X}^{\vec{m}_L}\bar{X}\right)P_{\vec{S}_1}. \tag{B.36}$$

Here, $C^1_{\vec{m}_L} = (\sin^2\alpha)^{|m_L|}(\cos^2\alpha)^{n-|m_L|}$ and $C^2_{\vec{m}_L} = (i\sin\alpha\cos\alpha)^n(-1)^{|m_L|}$.

The first term represents symmetric contributions and can be interpreted as incoherent errors. We write the first and the second terms as:

$$P_{\vec{S}_1}\left(\mathcal{E}(\rho_i)\right)P_{\vec{S}_1} = P_{\vec{S}_1}\left(\mathcal{N}_{\text{sym}}(\rho_i) + \mathcal{N}_{\text{asym}}(\rho_i)\right)P_{\vec{S}_1} \tag{B.37}$$

This implies that the full channel can be written as:

$$\left(\circ_{t=1}^{t_f}(M\circ\mathcal{E})\right)(\rho_i) = \left(\circ_{t=1}^{t_f}(M\circ(\mathcal{N}_{\text{sym}}+\mathcal{N}_{\text{asym}}))\right)(\rho_i) \tag{B.38}$$

Now we want to show that the full noise channel can be approximated by the local dephasing channel as:

$$\left(\circ_{t=1}^{t_f}(M\circ\mathcal{E})\right)(\rho_i) \approx \left(\circ_{t=1}^{t_f}(M\circ\mathcal{E}_p)\right)(\rho_i) \tag{B.39}$$

Let us first define the following density matrix:

$$\rho_t = \left(\circ_{t+1}^{t_f}(M\circ\mathcal{E})\right)\circ\left(\circ_{t=1}^{t}(M\circ\mathcal{N}_{\text{sym}})\right)(\rho_i) \tag{B.40}$$

This is the density matrix of the state evolved to time t, undergoing only the $\mathcal{N}_{\text{sym}}$ noise, and subsequently experiencing the $\mathcal{E}$ noise at each time step beyond t. Using a telescoping sum, we have:

$$\sum_{t=0}^{t=t_f} \rho_{t+1} - \rho_t = \left(\circ_{t=1}^{t_f}(M \circ \mathcal{N}_{\text{sym}})\right)(\rho_i) - \left(\circ_{t=1}^{t_f}(M \circ \mathcal{E})\right)(\rho_i) \tag{B.41}$$

This implies:

$$|| \left(\circ_{t=1}^{t_f}(M \circ \mathcal{N}_{\text{sym}})\right)(\rho_i) - \left(\circ_{t=1}^{t_f}(M \circ \mathcal{E})\right)(\rho_i)||_1 \leq \sum_{t=0}^{t=t_f} ||\rho_{t+1} - \rho_t||_1 \tag{B.42}$$

Now let's bound each $||\rho_{t+1} - \rho_t||_1$:

$$\rho_{t+1} - \rho_t = \left(\circ_{t+2}^{t_f}(M \circ \mathcal{E})\right) \circ (M \circ \mathcal{N}_{\text{sym}} - M \circ \mathcal{E}) \circ (\circ_{t=1}^{t}(M \circ \mathcal{N}_{\text{sym}}))(\rho_i) \tag{B.43}$$

Using the data processing inequality we have:

$$||\rho_{t+1} - \rho_t||_1 \leq ||(M \circ \mathcal{N}_{\text{sym}} - M \circ \mathcal{E})(\bar{\rho})||_1 = ||(M \circ \mathcal{N}_{\text{asym}})(\bar{\rho})||_1 \tag{B.44}$$

where $\bar{\rho} = (\circ_{t=1}^{t}(M \circ \mathcal{N}_{\text{sym}}))(\rho_i)$. Form the Eq. (B.36) we have:

$$(M \circ \mathcal{N}_{\text{sym}} - M \circ \mathcal{E})(\bar{\rho}) = \sum_{\vec{S}_1} \sum_{\vec{m}_L} C^2_{\vec{m}_L} P_{\vec{S}_1} \left(\hat{X}^{\vec{m}_L}(\bar{\rho})\hat{X}^{\vec{m}_L}\bar{X}\right) P_{\vec{S}_1} \otimes |\vec{S}_1\rangle\langle\vec{S}_1| \tag{B.45}$$

Bounding the right-hand side and again using data processing inequality:

$$|| \sum_{\vec{S}_1} \sum_{\vec{m}_L} C^2_{\vec{m}_L} P_{\vec{S}_1} \left(\hat{X}^{\vec{m}_L}(\bar{\rho})\hat{X}^{\vec{m}_L}\bar{X}\right) P_{\vec{S}_1} \otimes |\vec{S}_1\rangle\langle\vec{S}_1|| \leq \sum_{\vec{m}_L} |C^2_{\vec{m}_L}| \; || \left(\hat{X}^{\vec{m}_L}(\bar{\rho})\hat{X}^{\vec{m}_L}\bar{X}\right) || \tag{B.46}$$

Noting that $|| \left(\hat{X}^{\vec{m}_L}(\bar{\rho})\hat{X}^{\vec{m}_L}\bar{X}\right) ||_1 = ||\bar{\rho}||_1 = 1$, we get:

$$||\rho_{t+1} - \rho_t||_1 \leq \sum_{\vec{m}_L} |C^2_{\vec{m}_L}| = (2 \sin\alpha \cos\alpha)^n \tag{B.47}$$

This finally results in:

$$|| \left(\circ_{t=1}^{t_f}(M \circ \mathcal{E})\right)(\rho_i) - \left(\circ_{t=1}^{t_f}(M \circ \mathcal{N}_{\text{sym}})\right)(\rho_i)||_1 \leq (t_f + 1)\,(\sin(2\alpha))^n. \tag{B.48}$$

This term will vanish as long as $\sin(2\alpha) \leq 1$. Now, let us define the Pauli-twirled probability $p = \sin^2 \alpha$. Then, we have $C^1_{\vec{m}_L} = (p)^{|m_L|}(1-p)^{n-|m_L|}$, which results in:

$$\left(\circ_{t=1}^{t_f}(M \circ \mathcal{N}_{\text{sym}})\right)(\rho_i) = \left(\circ_{t=1}^{t_f}(M \circ \mathcal{E}_p)\right)(\rho_i) \tag{B.49}$$

where $\mathcal{E}_p$ is a dephasing channel with probability p. This finally establishes that:

$$|| \left(\circ_{t=1}^{t_f}(M \circ \mathcal{E})\right)(\rho_i) - \left(\circ_{t=1}^{t_f}(M \circ \mathcal{E}_p)\right)(\rho_i)||_1 \leq n\,(\sin(2\alpha))^n \tag{B.50}$$

Which is exponentially small as long as $\sin^2 \alpha = p \leq \frac{1}{2}$, thereby completing the proof.

B.2.3 Coherent error and forced measurements

Here, we establish an explicit mapping (via postselection) between the density matrix of the X-decohered noisy resource state and the corresponding syndrome distribution in the circuit model. This allows us to relate the CMI of the syndromes in a circuit subjected to coherent error to that of a resource state subjected to the corresponding coherent error.

The mapping that we introduce can be interpreted as a form of post-selection. We define a post-selection map $\mathcal{P}$ by

$$\mathcal{P}(\rho_\theta) = \frac{\Pi^C \rho_\theta \Pi^C}{\text{Tr}[\Pi^C \rho_\theta]}, \tag{B.51}$$

where Π^C projects all code qubits onto the $|+\rangle$ state. In words, we measure all of the code qubits of the resource state and only keep the result with $|+\rangle$ on every site.

In the context of the main text, we show that this post-selection yields the syndrome density matrix:

$$\mathcal{P}(\rho_\theta) = \bar{\rho}_\theta, \tag{B.52}$$

where ρ_θ is the density matrix of the X-decohered noisy resource state, and $\bar{\rho}_\theta$ represents the syndrome distribution in the corresponding circuit.

To validate Eq. (B.52) and establish the correspondence, we analyze the effect of forced measurements on a code qubit subject to noise. We demonstrate that such a measurement effectively transfers the error to the teleported qubit, thereby implementing the same noise process. Specifically, we consider the teleportation of qubit i from time t to time $t+1$, realized via measurements at time slices t and $t + \frac{1}{2}$, as described in Refs. [35, 88]. The

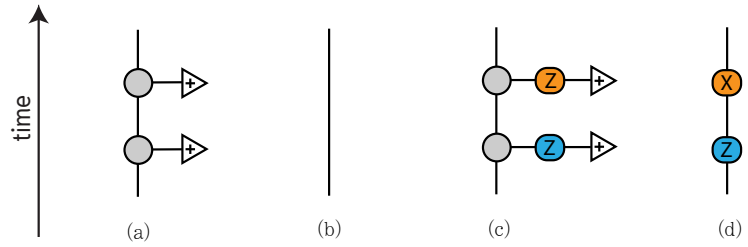


Figure B.1: (a) Code qubits at times t and $t + \frac{1}{2}$ are post-selected into the $|+\rangle$ state, indicated by contraction with a triangle symbol. (b) This post-selection implements perfect teleportation, effectively enacting a time evolution from time t to $t + 1$. (c) A coherent error unitary is applied to the code qubits, represented by the blue and orange gates. The post-selection is performed afterward. The Kraus operators of both noise channels are constructed using only Pauli Z. (d) The net effect is equivalent to applying the same coherent error during teleportation. In the circuit picture, the blue channel is applied directly to the teleported qubit, while the orange channel is mapped to an equivalent channel in which each Pauli Z operator is replaced by a Pauli X.

net effect is to propagate the quantum state to time $t+1$, with the error carried along the trajectory, as shown in Fig. B.1.

We do not provide a direct proof of the correspondence illustrated in Fig. B.1, as the argument is straightforward and involves applying gates and measurements on a three-qubit system. Given this correspondence, we have that post-selecting the code qubits precisely implements the circuit model, since the syndrome qubits are extracted identically in both the MBQC and circuit frameworks.

Having established this mapping, we now demonstrate that it preserves the CMI. In particular, this shows that a divergence of Markov length in one framework implies a corresponding divergence in the other. The proof proceeds in two steps. First, we show that the CMI associated with the detector cells is preserved under the mapping and coincides with that of the deterministic components in the circuit. Second, we use the argument from the main text, namely that the total CMI depends solely on the contributions from the detector cells, which guarantees the preservation of CMI across the two setups.

The post-measurement distribution of the noisy resource state can be expressed as a joint distribution over detector cells and random components. In the noiseless case, measurement outcomes are entirely random except for those associated with detector cells.

This is reflected in the expectation values:

$$\left\langle \prod_{i\in D} X_i \right\rangle = 1, \quad \left\langle \prod_{j\in D'} X_j \right\rangle = 0 \quad \text{for } D' \text{ not a detector cell.}$$

This can be shown for noiseless resource state by identifying a stabilizer of the resource state of the form $P_{j'} \prod Z_j$, where $P_{j'}$ is a Pauli operator supported outside D', and $\prod Z_j$ is the only component overlapping with D'. The resulting anticommutation between $\prod Z_j$ and $\prod_{j\in D'} X_j$ ensures that the expectation value $\langle \prod_{j\in D'} X_j \rangle$ vanishes in the resource state.

In the presence of coherent noise, the condition $\langle \prod_{j\in D'} X_j \rangle = 0$ remains valid for the cluster state, as well as any other resource state admitting a stabilizer of the form $P_{j'} \prod Z_j$, with support disjoint from D' except for the Z-type terms. This ensures that the random components retain their full randomness.

To make this precise, consider the ideal resource state ρ_0, stabilized by such an operator. When an onsite coherent Z-type error is applied, the resulting state ρ_θ is stabilized by the conjugated operator $O_{j'} \prod Z_j$, where $O_{j'}$ acts exclusively on qubits outside of D'. Importantly, the $\prod Z_j$ term remains unaffected by the error, and hence continues to anticommute with $\prod_{j\in D'} X_j$. As a result, the expectation value $\left\langle \prod_{j\in D'} X_j \right\rangle = 0$ still holds in the noisy state, confirming that measurement outcomes on non-detector cells remain fully random.

This decomposition into detector cells and random components allows us to express the state as

$$\rho_\theta = \mathcal{E}^x_{p=\frac{1}{2}}\left(\mathcal{E}^z_\theta(\rho_{\mathrm{RS}})\right) = \sum_{m,c,s} P_m \, \delta(c, s \equiv m) \, |c, s\rangle \langle c, s| \,, \tag{B.53}$$

where P_m denotes the probability of observing the detector outcome m. The state $|c, s\rangle$ labels the measurement outcomes of the code qubits (c) and syndrome qubits (s). The delta function $\delta(c, s \equiv m)$ enforces consistency between the code and syndrome configurations with the expected detector outcome m.

Applying post-selection via Π^C, we obtain

$$\bar{\rho}_\theta \propto \Pi^C \rho_\theta \, \Pi^C = \sum_{m,c,s} P_m \, \delta(c, s \equiv m) \Pi^C \, |c, s\rangle \langle c, s| \, \Pi^C \propto \sum_{s,m} P_m \, \delta(s \equiv m) \, |s\rangle \langle s| \,, \tag{B.54}$$

where $|s\rangle \propto \Pi^C |c, s\rangle$ are the normalized projections onto a fixed configuration of the code qubits, which we suppress for simplicity. Since each detector cell acts on both code and

syndrome qubits, fixing the code configuration ensures consistency with the corresponding syndrome outcomes, which is enforced by the constraint $\delta(s \equiv m)$.

On the other hand, the syndrome distribution in the circuit model can be decomposed as

$$\bar{\rho}_\theta = \sum_{s,d} P_d \, \delta(s \equiv d) \, |s\rangle \langle s| \,, \tag{B.55}$$

where d labels the deterministic components of the circuit.

By comparing the two expressions, we conclude that the probabilities are proportional. Given the one-to-one correspondence between detector cells in MBQC and deterministic components in the circuit, this proportionality implies equality:

$$P_{d_m} = P_m, \tag{B.56}$$

establishing that the probabilities associated with detector outcomes in MBQC match those of the corresponding deterministic components in the circuit model.

As discussed in the main text, the entropy of any region in both the circuit and the resource state can be decomposed into contributions from the deterministic (or detector) part and the random part. Since the random part does not contribute to the CMI, this equivalence of probabilities implies that the CMI is identical across the MBQC and circuit models, independent of the nature of the underlying noise.

B.3 Coherent error in MBQC

In the previous appendix, we were unable to find a simple form for the noise on the mixed state given coherent noise in the circuit model, except in the case of the repetition code. Here, we ask whether there is an associated mixed-state phase transition, given coherent noise acting on the MBQC resource state. We find that, indeed, coherent noise on the resource state maps to iid incoherent noise on the mixed state. This establishes a connection between the fault-tolerance threshold of MBQC to coherent noise and a mixed-state phase transition. In particular, we show that there exists a channel $\mathcal{G}$ that maps MBQC with coherent errors to MBQC with incoherent errors:

$$\mathcal{G}\left(\mathcal{E}^x_{p=1/2}\left(\mathcal{E}^z_\theta(\rho_{\mathrm{RS}})\right)\right) = \mathcal{E}^x_{p=1/2}\left(\mathcal{E}^z_p(\rho_{\mathrm{RS}})\right) \tag{B.57}$$

Where $\mathcal{E}^z_p$ is the Pauli-twirled of $\mathcal{E}^z_\theta$. This shows that the existence of a recovery channel for incoherent errors implies the existence of a recovery channel for coherent errors in MBQC.

To show the existence of $\mathcal{G}$, we begin by defining ρ_θ as the state obtained by applying coherent Z-errors, followed by maximal X-dephasing to the resource state:

$$\rho_\theta = \mathcal{E}^x_{p=1/2}\left(\mathcal{E}^z_\theta(\rho_{\mathrm{RS}})\right) \tag{B.58}$$

Here, $\mathcal{E}^z_\theta$ denotes a coherent rotation error, and $\mathcal{E}^x_{p=1/2}$ denotes a maximally dephasing X-error channel. To analyze the behavior under the feedback channel F, we introduce the measurement channel $\mathcal{E}^x_D$ representing the measurement of the detector cells. Note that

$$\mathcal{E}^x_{p=1/2} = \mathcal{E}^x_{p=1/2} \circ \mathcal{E}^x_D. \tag{B.59}$$

We define the intermediate state

$$\rho'_\theta = \mathcal{E}^x_D\left(\mathcal{E}^z_\theta(\rho_{\mathrm{RS}})\right), \tag{B.60}$$

so that $\rho_\theta = \mathcal{E}^x_{p=1/2}(\rho'_\theta)$.

Expanding the coherent error, we have:

$$\mathcal{E}^x_D\left(\mathcal{E}^z_\theta(\rho_{\mathrm{RS}})\right) = \sum_{\vec{D}} \sum_{\vec{m}} C_{\vec{m}} P_{\vec{D}} \hat{Z}^{\vec{m}_L} \rho_{\mathrm{RS}} \hat{Z}^{\vec{m}_R} P_{\vec{D}}. \tag{B.61}$$

The terms in this expansion vanish unless

$$\hat{Z}^{\vec{m}_L} \hat{Z}^{\vec{m}_R} P_{\vec{D}} = P_{\vec{D}} \hat{Z}^{\vec{m}_L} \hat{Z}^{\vec{m}_R}, \tag{B.62}$$

So we take $\hat{Z}^{\vec{m}_R} = \hat{Z}^{\vec{m}_{LR}} \hat{Z}^{\vec{m}_L}$.

$$\mathcal{E}^x_{p=1/2}\left(\mathcal{E}^z_\theta(\rho_{\mathrm{RS}})\right) = \mathcal{E}^x_{p=1/2}\left[\sum_{\vec{m}_L, \vec{m}_{LR}} C_{\vec{m}} \hat{Z}^{\vec{m}_L} \rho_{\mathrm{RS}} \hat{Z}^{\vec{m}_{LR}} \hat{Z}^{\vec{m}_L}\right], \tag{B.63}$$

We now want to study $F\left(\mathcal{E}^x_{p=1/2}(\mathcal{E}^z_\theta(\rho_{\mathrm{RS}}))\right)$. First, note that since $\hat{Z}^{\vec{m}_{LR}}$ commutes with the detector cells—which are all constructed as products of Pauli X operators—there must exist an operator $\hat{X}^{\vec{n}_{LR}}$ such that $\hat{Z}^{\vec{m}_{LR}} \hat{X}^{\vec{n}_{LR}}$ is a stabilizer of the resource state. In other words, the resource state is stabilized by $\hat{Z}^{\vec{m}_{LR}} \hat{X}^{\vec{n}_{LR}}$. We can therefore write:

$$F\left(\mathcal{E}^x_{p=1/2}(\mathcal{E}^z_\theta(\rho_{\mathrm{RS}}))\right) = \sum_{\vec{m}_L, \vec{n}_{LR}} C'_{\vec{m}} F\left(\mathcal{E}^x_{p=1/2}(\hat{Z}^{\vec{m}_L} \rho_{\mathrm{RS}} \hat{Z}^{\vec{m}_L} \hat{X}^{\vec{n}_{LR}})\right) \tag{B.64}$$

By introducing $\tilde{\rho}_{\vec{m}_L,\vec{C}}$ as the following:

$$\mathcal{E}^x_{p=1/2}(\hat{Z}^{\vec{m}_L}\rho_{\mathrm{RS}}\hat{Z}^{\vec{m}_L}) = \frac{1}{2^{|\vec{C}|}}\sum_{\vec{C}}|\vec{C}\rangle\langle\vec{C}|\otimes\tilde{\rho}_{\vec{m}_L,\vec{C}}, \tag{B.65}$$

We can write a term of Eq. (B.64) as:

$$\mathcal{E}^x_{p=1/2}(\hat{Z}^{\vec{m}_L}\rho_{\mathrm{RS}}\hat{Z}^{\vec{m}_L}\hat{X}^{\vec{n}_{LR}}) = \frac{1}{2^{|\vec{C}|}}\sum_{\vec{C}}(-1)^{f(\vec{C})}|\vec{C}\rangle\langle\vec{C}|\otimes\tilde{\rho}_{\vec{m}_L,\vec{C}}, \tag{B.66}$$

The factor $(-1)^{f(\vec{C})}$ in the sum corresponds to the sign of $\hat{X}^{\vec{n}_{LR}}$ after the measurements.

The feedback channel acts as:

$$F(\cdot) = \mathrm{Tr}_C\left(\sum_{\vec{C}}U_{\vec{C}}P_{\vec{C}}(\cdot)P_{\vec{C}}U^\dagger_{\vec{C}}\right), \tag{B.67}$$

so we obtain:

$$\begin{aligned} F\left(\mathcal{E}^x_{p=1/2}(\hat{Z}^{\vec{m}_L}\rho_{\mathrm{RS}}\hat{Z}^{\vec{m}_L}\hat{X}^{\vec{n}_{LR}})\right) &= \frac{1}{2^{|\vec{C}|}}\sum_{\vec{C}}(-1)^{f(\vec{C})}\,\mathrm{Tr}_C\left(|\vec{C}\rangle\langle\vec{C}|\right)\tilde{\rho}_{\vec{m}_L} \\ &= \frac{1}{2^{|\vec{C}|}}\left(\sum_{\vec{C}}(-1)^{f(\vec{C})}\right)\tilde{\rho}_{\vec{m}_L} = 0, \end{aligned} \tag{B.68}$$

Since $\hat{X}^{\vec{n}_{LR}}$ is not a detector cell itself, the sign can be either positive or negative depending on the measurement outcome. This results in vanishing of this term when summing over $\vec{C}$. Therefore, the feedback channel annihilates all non-symmetric terms.

We conclude:

$$F\left(\mathcal{E}^x_{p=1/2}\left(\mathcal{E}^z_\theta(\rho_{\mathrm{RS}})\right)\right) = F\left(\mathcal{E}^x_{p=1/2}\left(\mathcal{E}^z_p(\rho_{\mathrm{RS}})\right)\right), \tag{B.69}$$

for $p = \sin^2\alpha$, establishing that coherent errors become effectively incoherent under the feedback channel. Therefore we have:

$$\mathcal{G} = F^{-1}\circ F \tag{B.70}$$

Appendix C

Supplementary Material for Chapter 4

C.1 Punctured coherent information and local recoverability

This Appendix proves Theorem 4.1. Fix a family of topological stabilizer codes on a D-dimensional lattice with physical register $Q = Q(L)$ and distance $d(L)$, and let $\rho^{(L)}_{RQ,\mathrm{noisy}}$ and $\rho^{(L)}_{RQ,\mathrm{ideal}}$ denote the corresponding noisy and target reference–output states appearing in Theorem 4.1. Throughout we assume the hypotheses of Theorem 4.1, namely: (i) the mutual informations match asymptotically,

$$\left| I(R{:}Q)_{\rho^{(L)}_{RQ,\mathrm{noisy}}} - I(R{:}Q)_{\rho^{(L)}_{RQ,\mathrm{ideal}}} \right| \xrightarrow[L\to\infty]{} 0, \tag{C.1}$$

(ii) there exists a family of recovery channels $\{\mathcal{R}_L\}_L$ implementable by a quasi-local circuit such that

$$\left\| (\mathrm{id}_R \otimes \mathcal{R}_L)\big(\rho^{(L)}_{RQ,\mathrm{noisy}}\big) - \rho^{(L)}_{RQ,\mathrm{ideal}} \right\|_1 \le \varepsilon_L \xrightarrow[L\to\infty]{} 0, \tag{C.2}$$

and (iii) the code distance obeys $d(L) \ge cL^{\alpha}$ for some constants $c, \alpha > 0$ and all sufficiently large L.

The proof has two components. First, quasi-locality of $\mathcal{R}_L$ induces a bounded light cone: puncturing a region $Q_p(L)$ can only affect the action of $\mathcal{R}_L$ inside an enlarged neighborhood $Q_{\mathrm{LC}}(L)$. Second, using the distance lower bound we choose $Q_p(L)$ so that $Q_{\mathrm{LC}}(L)$ is correctable for the code. This yields a recovery from the punctured noisy state with error controlled by ε_L (Lemma C.1). Finally, a simple data-processing argument implies that puncturing does not change the mutual information asymptotically, proving Theorem 4.1.

C.1.1 Circuit locality and light cones

We recall the circuit notion of quasi-locality used in the main text. Let $(Q(L), d_{\mathrm{lat}})$ be the lattice metric space of physical qubits. A CPTP map $\mathcal{R}_L : Q \to Q$ is a *local channel circuit* of depth t_L and range w_L if it can be written as

$$\mathcal{R}_L = \mathcal{E}_L^{(t_L)} \circ \cdots \circ \mathcal{E}_L^{(1)}, \tag{C.3}$$

where each layer factorizes over disjoint supports,

$$\mathcal{E}_L^{(s)} = \bigotimes_j \mathcal{E}_{L,j}^{(s)}, \qquad \operatorname{diam}\big(\operatorname{supp}(\mathcal{E}_{L,j}^{(s)})\big) \le w_L. \tag{C.4}$$

We say the circuit is *strictly local* if $t_L = O(1)$ and $w_L = O(1)$, and *quasi-local* if $t_L = \mathrm{polylog}(L)$ and $w_L = \mathrm{polylog}(L)$.

A depth-t_L, range-w_L circuit has a light-cone radius

$$r_L := c_0\, t_L w_L, \tag{C.5}$$

for a constant $c_0 > 0$ depending only on geometric conventions; in the quasi-local case $r_L = \mathrm{polylog}(L)$. Fix a puncture $Q_p(L) \subset Q$ and define the associated light-cone neighborhood

$$Q_{\mathrm{LC}}(L) := \mathsf{N}_{r_L}\big(Q_p(L)\big), \qquad Q_{\mathrm{LC}}^c(L) := Q \setminus Q_{\mathrm{LC}}(L), \tag{C.6}$$

where $\mathsf{N}_r(\cdot)$ is the metric r-neighborhood.

C.1.2 Proof of Theorem 4.1

Theorem 4.1 follows from the next lemma, which upgrades recovery of the unpunctured noisy state to recovery of the punctured noisy state once the enlarged puncture $Q_{\mathrm{LC}}(L)$ is correctable. Its proof is given in Sec. C.1.4.

Lemma C.1 (Punctured recovery)**.** *Assume correctability of the enlarged puncture, i.e. there exists a CPTP map $\mathcal{T}_L : Q_{\mathrm{LC}}^c(L) \to Q(L)$ such that*

$$(\mathrm{id}_R \otimes \mathcal{T}_L)\Big[(\mathrm{id}_R \otimes \mathrm{tr}_{Q_{\mathrm{LC}}})\big(\rho_{RQ,\mathrm{ideal}}^{(L)}\big)\Big] = \rho_{RQ,\mathrm{ideal}}^{(L)}. \tag{C.7}$$

Let $\rho_{RQ_{p^c}}^{(L)} := \mathrm{tr}_{Q_p(L)}\, \rho_{RQ,\mathrm{noisy}}^{(L)}$ be the punctured noisy state. Then there exists a recovery channel $\widetilde{\mathcal{R}}_L$ acting on $Q_{p^c}(L)$ such that

$$\big\|(\mathrm{id}_R \otimes \widetilde{\mathcal{R}}_L)\big(\rho_{RQ_{p^c}}^{(L)}\big) - \rho_{RQ,\mathrm{ideal}}^{(L)}\big\|_1 \le \varepsilon_L. \tag{C.8}$$

Moreover, $\widetilde{\mathcal{R}}_L$ may be chosen as

$$\widetilde{\mathcal{R}}_L = \mathcal{T}_L \circ \mathrm{tr}_{Q_{\mathrm{LC}}} \circ \mathcal{R}_L \circ \left(\tau_{Q_p(L)} \otimes \mathrm{id}_{Q_{p^c}(L)}\right). \tag{C.9}$$

i.e., append fresh qubits on the puncture in the state $\tau_{Q_p(L)}$, apply the original recovery $\mathcal{R}_L$ on Q, discard the causal-cone region $Q_{\mathrm{LC}}(L)$, and finally apply the reconstruction $\mathcal{T}_L$.

We now complete the proof of Theorem 4.1. By Lemma C.1, the recovered punctured state $(\mathrm{id}_R \otimes \widetilde{\mathcal{R}}_L)(\rho^{(L)}_{RQ_{p^c}})$ is ε_L-close in trace norm to $\rho^{(L)}_{RQ,\mathrm{ideal}}$. Since the reference system R has fixed dimension, continuity of mutual information implies

$$I(R{:}Q)_{(\mathrm{id}_R\otimes\widetilde{\mathcal{R}}_L)(\rho^{(L)}_{RQ_{p^c}})} - I(R{:}Q)_{\rho^{(L)}_{RQ,\mathrm{ideal}}} \xrightarrow[L\to\infty]{} 0. \tag{C.10}$$

On the other hand, by the data-processing inequality applied to the channel $\widetilde{\mathcal{R}}_L$ acting on $Q_{p^c}(L)$,

$$I(R{:}Q)_{(\mathrm{id}_R\otimes\widetilde{\mathcal{R}}_L)(\rho^{(L)}_{RQ_{p^c}})} \le I(R{:}Q_{p^c}(L))_{\rho^{(L)}_{RQ,\mathrm{noisy}}}. \tag{C.11}$$

Applying data processing again to the partial trace $\mathrm{tr}_{Q_p(L)}$ (a channel on Q) yields

$$I(R{:}Q_{p^c}(L))_{\rho^{(L)}_{RQ,\mathrm{noisy}}} \le I(R{:}Q)_{\rho^{(L)}_{RQ,\mathrm{noisy}}}. \tag{C.12}$$

Combining (C.10)–(C.12) with the mutual-information matching condition (C.1) gives:

$$I(R{:}Q_{p^c}(L))_{\rho^{(L)}_{RQ,\mathrm{noisy}}} - I(R{:}Q)_{\rho^{(L)}_{RQ,\mathrm{noisy}}} \xrightarrow[L\to\infty]{} 0, \tag{C.13}$$

which is the claim (4.24) of Theorem 4.1.

We now choose a puncture family $\{Q_p(L)\}_L$ with $|Q_p(L)|$ growing polynomially in L such that the enlarged neighborhood $Q_{\mathrm{LC}}(L)$ is correctable.

Fix a lattice-dependent constant $\kappa > 0$ such that any hypercubic region of linear size u contains at most κu^D qubits. Choose $Q_p(L) \subseteq Q(L)$ to be a bulk hypercubic region of linear size ℓ_L. Since $Q_{\mathrm{LC}}(L) = \mathsf{N}_{r_L}(Q_p(L))$ is contained in a hypercube of linear size $\ell_L + 2r_L$, we have

$$|Q_{\mathrm{LC}}(L)| \le \kappa\,(\ell_L + 2r_L)^D. \tag{C.14}$$

Because $r_L = \mathrm{polylog}(L)$ and $d(L) \ge cL^\alpha$, we choose ℓ_L so that $\kappa(2\ell_L)^D = d(L)$, which implies $r_L = o(\ell_L)$. In particular, for all sufficiently large L,

$$|Q_{\mathrm{LC}}(L)| < d(L). \tag{C.15}$$

Thus $Q_{\mathrm{LC}}(L)$ is a correctable region for the code family, and hence there exists a reconstruction map $\mathcal{T}_L$ satisfying (C.7). This establishes the hypotheses of Lemma C.1 and completes the proof of Theorem 4.1.

C.1.3 Causal-cone insensitivity

We formalize the causal-cone statement used implicitly in Lemma C.1.

Lemma C.2 (Light-cone causality for channel circuits)**.** *Let $\mathcal{R}_L$ be a local channel circuit of depth t_L and range w_L, and let r_L be as in Eq. (C.5). Then for any CPTP map Φ supported on $Q_p(L)$ and any state ρ_{RQ},*

$$\begin{aligned}&(\mathrm{id}_R \otimes \mathrm{tr}_{Q_{\mathrm{LC}}}) \circ (\mathrm{id}_R \otimes \mathcal{R}_L)(\rho_{RQ}) \\ &= (\mathrm{id}_R \otimes \mathrm{tr}_{Q_{\mathrm{LC}}}) \circ (\mathrm{id}_R \otimes \mathcal{R}_L) \circ (\mathrm{id}_R \otimes \Phi)(\rho_{RQ}).\end{aligned} \tag{C.16}$$

Proof. Let O be any observable supported on $R \cup Q^c_{\mathrm{LC}}(L)$. Consider the Heisenberg-picture evolution of O under $\mathrm{id}_R \otimes \mathcal{R}_L^\dagger$. Because each layer $\mathcal{E}_L^{(s)}$ factorizes into disjoint gates of diameter at most w_L, the support of O can expand by at most $O(w_L)$ in lattice diameter per layer. After t_L layers, the support of the backwards-evolved observable remains disjoint from $Q_p(L)$ by the definition $Q_{\mathrm{LC}}(L) = \mathsf{N}_{r_L}(Q_p(L))$ with $r_L = c_0 t_L w_L$. Hence

$$\begin{aligned}&\mathrm{tr}[O\,(\mathrm{id}_R \otimes \mathcal{R}_L)(\rho_{RQ})] \\ &= \mathrm{tr}\big[O\,(\mathrm{id}_R \otimes \mathcal{R}_L)\big((\mathrm{id}_R \otimes \Phi)(\rho_{RQ})\big)\big]\,.\end{aligned} \tag{C.17}$$

for all such O. Since equality of all expectation values on $R \cup Q^c_{\mathrm{LC}}(L)$ is equivalent to equality of the reduced states on $RQ^c_{\mathrm{LC}}(L)$, this proves Eq. (C.16). □

C.1.4 Proof of Lemma C.1

Proof. Fix any product state $\tau_{Q_p(L)}$ on $Q_p(L)$ and define the "pad" map

$$\mathcal{A}_L := \tau_{Q_p(L)} \otimes \mathrm{id}_{Q_{p^c}(L)}, \tag{C.18}$$

which maps states on $Q_{p^c}(L)$ to states on $Q(L)$ by adjoining fresh qubits in the puncture. Define the punctured-and-padded state

$$\widetilde{\rho}^{(L)}_{RQ} := (\mathrm{id}_R \otimes \mathcal{A}_L)\big(\rho^{(L)}_{RQ_{p^c}}\big). \tag{C.19}$$

Step 1: invariance outside the causal cone. Apply Lemma C.2 with Φ equal to the CPTP map $\mathrm{tr}_{Q_p(L)}$ followed by padding with $\tau_{Q_p(L)}$. Equivalently, puncturing-and-padding acts only on $Q_p(L)$, so it cannot influence the reduced recovered state outside $Q_{\mathrm{LC}}(L)$. Concretely,

$$\begin{aligned}&(\mathrm{id}_R \otimes \mathrm{tr}_{Q_{\mathrm{LC}}})\big[(\mathrm{id}_R \otimes \mathcal{R}_L)\big(\rho^{(L)}_{RQ,\mathrm{noisy}}\big)\big] \\ &= (\mathrm{id}_R \otimes \mathrm{tr}_{Q_{\mathrm{LC}}})\big[(\mathrm{id}_R \otimes \mathcal{R}_L)\big(\widetilde{\rho}^{(L)}_{RQ}\big)\big].\end{aligned} \tag{C.20}$$

Step 2: transfer the unpunctured recovery guarantee to the punctured-and-padded input. By contractivity of trace distance under partial trace, Eq. (C.2) implies

$$\Big\| (\mathrm{id}_R \otimes \mathrm{tr}_{Q_{\mathrm{LC}}}) \big[(\mathrm{id}_R \otimes \mathcal{R}_L)(\rho^{(L)}_{RQ,\mathrm{noisy}}) \big] - (\mathrm{id}_R \otimes \mathrm{tr}_{Q_{\mathrm{LC}}})(\rho^{(L)}_{RQ,\mathrm{ideal}}) \Big\|_1 \leq \varepsilon_L. \tag{C.21}$$

Using Eq. (C.20), we may replace the first term by the punctured-and-padded input:

$$\Big\| (\mathrm{id}_R \otimes \mathrm{tr}_{Q_{\mathrm{LC}}}) \big[(\mathrm{id}_R \otimes \mathcal{R}_L)(\widetilde{\rho}^{(L)}_{RQ}) \big] - (\mathrm{id}_R \otimes \mathrm{tr}_{Q_{\mathrm{LC}}})(\rho^{(L)}_{RQ,\mathrm{ideal}}) \Big\|_1 \leq \varepsilon_L. \tag{C.22}$$

Step 3: reconstruct the discarded causal cone using correctability. Apply the CPTP reconstruction $\mathcal{T}_L$ (assumed to exist by Eq. (C.7)) to both states in Eq. (C.22). By contractivity under CPTP maps,

$$\Big\| (\mathrm{id}_R \otimes \mathcal{T}_L)(\mathrm{id}_R \otimes \mathrm{tr}_{Q_{\mathrm{LC}}})(\mathrm{id}_R \otimes \mathcal{R}_L)(\widetilde{\rho}^{(L)}_{RQ}) - (\mathrm{id}_R \otimes \mathcal{T}_L)(\mathrm{id}_R \otimes \mathrm{tr}_{Q_{\mathrm{LC}}})(\rho^{(L)}_{RQ,\mathrm{ideal}}) \Big\|_1 \leq \varepsilon_L. \tag{C.23}$$

The second term equals $\rho^{(L)}_{RQ,\mathrm{ideal}}$ by Eq. (C.7), hence

$$\Big\| (\mathrm{id}_R \otimes \mathcal{T}_L)(\mathrm{id}_R \otimes \mathrm{tr}_{Q_{\mathrm{LC}}})(\mathrm{id}_R \otimes \mathcal{R}_L)(\widetilde{\rho}^{(L)}_{RQ}) - \rho^{(L)}_{RQ,\mathrm{ideal}} \Big\|_1 \leq \varepsilon_L. \tag{C.24}$$

Step 4: define the punctured recovery map. Define a recovery channel on $Q_{p^c}(L)$ by

$$\widetilde{\mathcal{R}}_L := \mathcal{T}_L \circ \mathrm{tr}_{Q_{\mathrm{LC}}} \circ \mathcal{R}_L \circ \mathcal{A}_L. \tag{C.25}$$

This has the explicit implementation in Eq. (C.9). Applying it to $\rho^{(L)}_{RQ_{p^c}}$ gives the left-hand side of Eq. (C.24), which proves Eq. (C.8). □

C.2 Apparent critical behaviour at $q = 0$

In the mixed state derived from the heralded dephasing channel, along the $q = 0$ line, the CMI and the input–output mutual information both show a sharp crossover at a value $p^* < 1$ that appears nearly independent of system size over the range we simulate. In particular, for $L = 21$–161 the curves in Fig. C.1 exhibit an apparent crossing near $p^* \simeq 0.7$, which one might naively interpret as a genuine classical-memory transition at $q = 0$.

However, the loop-model description of Nahum *et al.* [160] implies that along $q = 0$ there is, strictly speaking, no transition out of the short loop phase except at the extremal point $p = 1$, which we refer to as a perfect "metal". The states at $q = 0$ are always localized for $p < 1$, but the associated correlation length $\xi(p)$, which also sets the typical loop size, becomes enormous as one approaches the perfect metal. The sigma-model analysis gives

$$\xi(p) \sim (1-p)^{-2} \exp\left[\frac{c}{1-p}\right], \tag{C.26}$$

for some positive, non-universal constant c [160]. Thus the characteristic length scale grows essentially exponentially fast as $p \to 1$, even though the phase remains localized for any $p < 1$.

For a finite system of linear size L, the crossover away from the perfect-metal fixed point at $(p = 1, q = 0)$ occurs when L becomes comparable to $\xi(p)$. This defines an apparent critical point $p^*(L)$ along $q = 0$ via

$$L \sim \xi(p^*(L)) \,. \tag{C.27}$$

Because of the exponential dependence in Eq. (C.26), $p^*(L)$ drifts towards $p = 1$ only extremely slowly as L increases. Over any modest range of system sizes, $p^*(L)$ therefore appears almost L-independent and mimics a true critical point, exactly as seen in Fig. C.1.

To make this more quantitative, consider two crossings $p^*(L_0)$ and $p^*(L_1)$ at system sizes L_0 and L_1, with

$$L_i \sim \xi(p^*(L_i)) \,, \qquad i = 0, 1. \tag{C.28}$$

Using Eq. (C.26) and eliminating the non-universal prefactor, we obtain:

$$\ln \frac{L_1}{L_0} \simeq -2 \ln\left(\frac{1-p^*(L_1)}{1-p^*(L_0)}\right) + c\left(\frac{1}{1-p^*(L_1)} - \frac{1}{1-p^*(L_0)}\right). \tag{C.29}$$

As a concrete estimate, suppose that for $L_0 \simeq 150$ we see an apparent crossing at $p^*(L_0) \simeq 0.70$, as in Fig. C.1, and we ask how large L_1 must be for this apparent crossing

to drift to $p^*(L_1) \simeq 0.75$. Substituting $p^*(L_0) = 0.70$ and $p^*(L_1) = 0.75$ into Eq. (C.29) yields

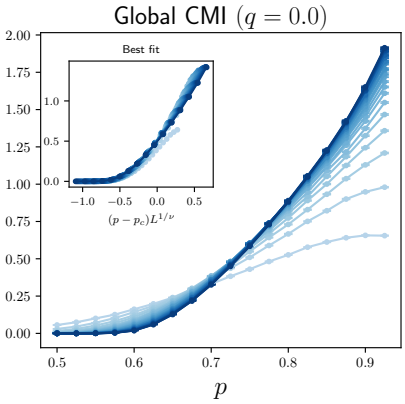
$$\ln \frac{L_1}{L_0} \simeq -2 \ln\left(\frac{1-0.75}{1-0.70}\right) + c\left(\frac{1}{1-0.75} - \frac{1}{1-0.70}\right)$$
$$\approx 0.36 + 0.67\, c. \tag{C.30}$$

The constant c is non-universal and depends on microscopic details; to get an order-of-magnitude estimate we may take $c \sim 1$, in which case

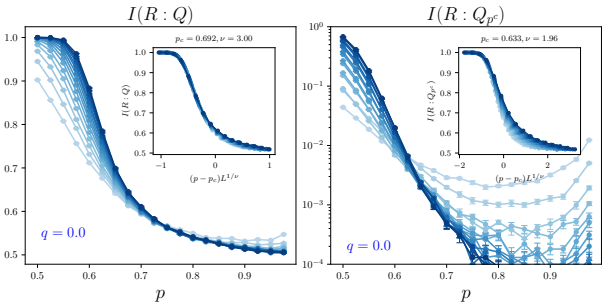
$$\ln \frac{L_1}{L_0} \approx 1.0 \qquad \Rightarrow \qquad \frac{L_1}{L_0} \sim e^{1.0} \approx 2.8. \tag{C.31}$$

Thus, in order for the apparent crossing to drift from $p^* \approx 0.70$ to $p^* \approx 0.75$ we would need to increase the system size by roughly a factor of three. Starting from $L_0 \approx 150$, this corresponds to $L_1 \sim 400$. Larger values of c would only increase the required system size. This back-of-the-envelope estimate explains why, over the accessible range $L \lesssim 160$, the apparent critical point along $q = 0$ in Fig. C.1 appears essentially pinned near $p^* \approx 0.7$.

In summary, the numerical results at $q = 0$ show an apparent critical behavior—a sharp, almost size-independent crossing in CMI and input–output mutual information—but the loop-model analysis indicates that this is a finite-size crossover controlled by the perfect-metal fixed point at $p = 1$, not a true critical phase. A detailed finite-size scaling of the input–output mutual information (Fig. C.1b) is consistent with a single localized phase along $q = 0$ with an exponentially large correlation length near $p = 1$.



(a) The global CMI as a function of p for fixed $q = 0.0$, and for different sizes of the surface code patch $L \times L$, with $L = 21$–161. There is an apparent crossing at $p^* \to 0.7$ that drifts only very weakly with L. The traces are averaged over 5×10^4 samples.



(b) The input–output mutual information of the encoding channel also shows an apparent crossing as a function of p at fixed $q = 0.0$. However, the punctured version shows a crossing at a significantly different p_c, suggesting that this transition is a finite-size effect, with strong drift in the scaling curves as L is varied.

Figure C.1: Apparent critical behaviour at $q = 0.0$. Along the boundary of the CPLC phase diagram, the CMI and input–output mutual information exhibit an apparent crossing at $p^* < 1$, which we interpret as a finite-size crossover controlled by the perfect-metal fixed point at $p = 1$ rather than a true phase transition.

Appendix D

Supplementary Material for Chapter 5

D.1 Emergence of canonical ensemble

Consider a Hamiltonian H with a symmetry described by a group G, where H commutes with the unitary representations $U(g)$ for all $g \in G$. Consider a one dimensional irrep Λ of the symmetry, labeled by the eigenvalues of the group generators, $U(g)\left|\Lambda,\alpha\right\rangle = \Lambda(g)\left|\Lambda,\alpha\right\rangle$, with α labeling the multiple representations of the irrep in the local tensor product Hilbert space. Note that the requirement of a one dimensional irrep is naturally satisfied for an Abelian group.

The projector onto the irrep space of Λ is given by

$$\Pi_\Lambda = \sum_\alpha \left|\Lambda,\alpha\right\rangle\left\langle\Lambda,\alpha\right| .$$

A strongly symmetric mixed state ρ in the symmetric subspace is one that satisfies the constraint

$$\rho = \Pi_\Lambda \rho \Pi_\Lambda .$$

A strongly symmetric quantum channel $\mathcal{E} = \sum_a K_a(\cdot)K_a^\dagger$ is one whose Kraus operators $\{K_a\}$ commute with the symmetry: $\forall g \in G, [K_a, U(g)] = 0$ [38, 7, 6]. Importantly, a strongly symmetric quantum channel preserved the strong symmetry:

Lemma D.1. *A strongly symmetric mixed state remains in the symmetric subspace when evolved with a strongly symmetric quantum channel.*

Proof. Suppose the initial state ρ_i is in the symmetric subspace, i.e., $\rho_i = \Pi_\Lambda \rho_i \Pi_\Lambda$. Then, the evolved state $\rho_f = \mathcal{E}(\rho_0)$ under a strongly symmetric channel $\mathcal{E} = \sum_a K_a(\cdot)K_a^\dagger$ also satisfies the condition:

$$\Pi_\Lambda \rho_f \Pi_\Lambda = \sum_a \Pi_\Lambda K_a \rho_0 K_a^\dagger \Pi_\Lambda = \sum_a K_a \Pi_\Lambda \rho_0 \Pi_\Lambda K_a^\dagger = \sum_a K_a \rho_0 K_a^\dagger = \rho_f. \tag{D.1}$$

□

Now, we want to show that the canonical ensemble is a state that minimizes the free energy associated with a local Hamiltonian when the states are restricted to a symmetric subspace. This would imply that such a state would naturally be the steady state of a symmetric mixing quantum process in the presence of a bath.

The free energy associated with a state at a temperature β^{-1} is given by,

$$F(\rho) = \operatorname{Tr} \rho H - \beta^{-1} S(\rho). \tag{D.2}$$

We can prove the following result,

Lemma D.2. *The canonical ensemble minimizes the free energy for all mixed states in the symmetric subspace. Specifically, for any state ρ such that $\rho = \Pi_\Lambda \rho \Pi_\Lambda$, we have,*

$$F(\rho_{\beta,\Lambda}) - F(\rho) \leq 0. \tag{D.3}$$

Proof. We study the quantity,

$$\operatorname{Tr} \rho \ln \rho_{\beta,\Lambda} = \operatorname{Tr} \rho \ln \Pi_\Lambda - \beta \operatorname{Tr} \rho H - \ln \operatorname{Tr} \Pi_\Lambda e^{-\beta H}. \tag{D.4}$$

For strongly symmetric states $\rho = \Pi_\Lambda \rho \Pi_\Lambda$, the first term above, $\operatorname{Tr} \rho \ln \Pi_\Lambda$, is zero, since $\ln \Pi_\Lambda$ is zero on the support of ρ.

For the canonical ensemble, we have,

$$F(\rho_{\beta,\Lambda}) = \operatorname{Tr} \rho_{\beta,\Lambda} H + \beta^{-1} \operatorname{Tr} \rho_{\beta,\Lambda} \ln \rho_{\beta,\Lambda} = -\beta^{-1} \ln \operatorname{Tr} \Pi_\Lambda e^{-\beta H}. \tag{D.5}$$

For any strongly symmetric state ρ, we have,

$$\begin{aligned} \operatorname{Tr} \rho \ln \rho_{\beta,\Lambda} &= -\beta \operatorname{Tr} \rho H - \ln \operatorname{Tr} \Pi_\Lambda e^{-\beta H} \\ &= -\beta (F(\rho) - \beta^{-1} \operatorname{Tr} \rho \ln \rho) - \ln \operatorname{Tr} \Pi_\Lambda e^{-\beta H} \\ &= \beta \left(F(\rho_{\beta,\Lambda}) - F(\rho) \right) + \operatorname{Tr} \rho \ln \rho. \end{aligned} \tag{D.6}$$

Thus,

$$\begin{aligned} F(\rho_{\beta,\Lambda}) - F(\rho) &= \beta^{-1} \left(\operatorname{Tr} \rho \ln \rho_{\beta,\Lambda} - \operatorname{Tr} \rho \ln \rho \right) \\ &= -\beta^{-1} D(\rho || \rho_{\beta,\Lambda}). \end{aligned} \tag{D.7}$$

By the positivity of relative entropy, we have the desired result. □

D.2 Two qubit example with $\mathbb{Z}_2$ symmetry

In this section, we pedagogically examine the case of two qubits with the $\mathbb{Z}_2$ symmetry generated by X_1X_2. The persistence of entanglement results of Theorem 5.2 are valid for such symmetry action, but some others are not due to the small system size. More specifically, the projection term Π_λ cannot be removed from the EC $[U_A(g), H] \neq 0$.

Indeed, consider the 2-qubit XYZ Hamiltonian

$$H = X_1X_2 + J\left(\frac{1+\gamma}{2}Y_1Y_2 + \frac{1-\gamma}{2}Z_1Z_2\right) = X_1X_2 + JY_1Y_2(\gamma\Pi_{+1} + \Pi_{-1}),$$

with added restriction to the $\mathbb{Z}_2$-symmetric sector $X_1X_2 = +1$. The negativity of its canonical ensemble state $\rho_{\beta,+1} \propto e^{-\beta H}\Pi_{+1}$ can be computed to be $\mathcal{N}(\rho_{\beta,+1}) = \frac{1}{2}|\tanh(\beta J\gamma)|$. Thus, for all temperatures, $\rho_{\beta,+1}$ is separable if $\gamma = 0$. However, if $J \neq 0$, H does not commute with the partial symmetry action and the SEC is satisfied:

$$[X_1, H] = iJ(Z_1Y_2 - Y_1Z_2) \neq 0,$$

implying persistence of symmetric entanglement. If one instead includes the projection Π_{+1}, then we have $[X_1, H\Pi_{+1}] = 2iJ\gamma Z_1Y_2\Pi_{+1} = 0$, as expected from the separability of $\rho_{\beta,+1}$. One might notice that this counterexample relies on the projections $\Pi_{\pm 1}$ directly appearing in the expression of the Hamiltonian, which should not happen for a system of many parties and a local Hamiltonian.

We can visually represent this example in the case of $\gamma = 0$ with the diagram in Fig. D.1. It is an accurate representation of a slice in the space of quantum states akin to Fig. 5.1. We can see that the Gibbs state (black dashed line, $J = 1$) becomes symmetrically entangled as soon as $\beta > 0$, but has sudden death of (non-symmetric) entanglement at low enough β. At the same time, the canonical ensemble (thick line from $\frac{1}{2}\Pi_-$ to $|\Psi_-\rangle$) becomes entangled as soon as $\beta > 0$.

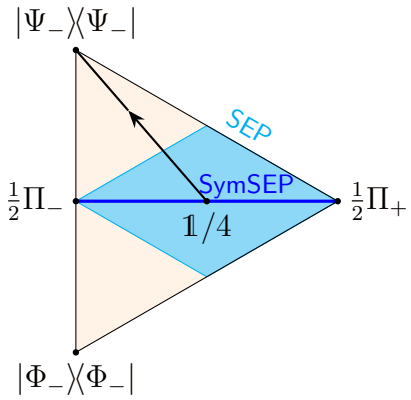


Figure D.1: Subset of quantum states of two qubits formed by convex mixtures of the Bell states $|\Psi_-\rangle \propto |01\rangle - |10\rangle$ and $|\Phi_-\rangle = |00\rangle - |11\rangle$, and the maximally mixed $X_1X_2 = +1$ state $\frac{1}{2}\Pi_+ = \frac{1}{2}|00\rangle\langle 00| + \frac{1}{2}|11\rangle\langle 11|$. The dashed line traces the path of the Heisenberg antiferromagnetic thermal states $e^{-\beta \mathbf{S}_1\cdot\mathbf{S}_2}$, or, equivalently, of the Werner states $\lambda|\Psi_-\rangle\langle\Psi_-| + (1-\lambda)\frac{1}{4}\mathbb{1}$. Meanwhile, the thick black line going from $\frac{1}{2}\Pi_-$ to $|\Psi_-\rangle$ traces the path of the canonical ensemble $e^{-\beta \mathbf{S}_1\cdot\mathbf{S}_2}\Pi_{-1}$ in the $X_1X_2 = -1$ sector. The SEP region in cyan consists of separable states, with the blue SymSEP line inside it indicating the symmetrically separable ones. Their loci were found via the Peres-Horodecki criterion, since it is necessary and sufficient for two-qubit systems [171, 94]. This illustration was inspired by Fig. 16.8(b) of [19].

D.3 Symmetric entanglement

We briefly recall the definition of separability and symmetric separability. A state ρ is separable for a bipartition $A|B$ if it can be represented as a convex mixture of tensor product states, $\rho = \sum_i p_i \rho_{A,i} \otimes \rho_{B,i}$. A separable state ρ is also symmetrically separable if all $\rho_{A,i} \otimes \rho_{B,i}$ are also (weakly) symmetric.

Following these definitions, and the definitions of weak and strong symmetries, we prove the following lemmas,

Lemma D.3. *A strongly symmetric state ρ is symmetrically separable if, and only if, it is separable.*

Proof. ($\Rightarrow$) is true for any state. For the ($\Leftarrow$) direction, given a separable and strongly symmetric state $\rho = \sum_i p_i \rho_{A,i} \otimes \rho_{B,i}$, each tensor product state $\rho_{A,i} \otimes \rho_{B,i}$ is also strongly

symmetric due to the inheritance property of the strong symmetry [134]. In particular, they are also weakly symmetric. □

Lemma D.4. *If the symmetry action $U : G \to U(\mathcal{H})$ is on-site, $U = U_A \otimes U_B$, then a weakly symmetric state ρ is symmetrically separable if, and only if, each of its strongly symmetric parts $\rho_\lambda \propto \rho\Pi_\Lambda$ is separable.*

Proof. ($\Rightarrow$) Since the irrep projector Π_Λ decomposes as $\Pi_\Lambda = \sum_\lambda \Pi^A_\lambda \Pi^B_{\Lambda\overline{\lambda}}$, then if $\rho = \sum_i p_i \rho_{A,i} \otimes \rho_{B,i}$, we have the following separable decomposition of ρ_λ:

$$\rho_\Lambda \propto \rho\Pi_\Lambda = \sum_{i,\lambda} p_i (\rho_{A,i}\Pi^A_\lambda) \otimes (\rho_{B,i}\Pi^B_{\Lambda\overline{\lambda}}). \tag{D.8}$$

($\Leftarrow$) follows from ρ being a convex mixture of its strongly symmetric parts ρ_λ, each of which is separable, and thus symmetrically separable by Lemma D.3. □

These two lemmas prove our assertion in the main text: ρ is symmetrically entangled (i.e. *not* symmetrically separable) if, and only if, its strongly symmetric component ρ_Λ is also entangled for some charge Λ.

D.4 Local indistinguishability of the canonical and Gibbs ensembles

In this section, we explore whether the canonical ensemble $\rho_{\beta,\Lambda}$ is locally indistinguishable from the Gibbs ensemble ρ_β, up to an exponentially small error. We establish this for $\beta = 0$ for finite Abelian groups — which we use to develop intuition about the equivalence of EC and SEC in the main text — and for small enough $\beta < \beta_*$ if H is a k-local Hamiltonian.

We start by collecting some known facts about the symmetries and their representations. From the symmetry action of G on the Hilbert space $\mathcal{H}$, it gets decomposed into $\mathcal{H} = \bigoplus_\lambda V_\lambda^{\otimes m_\lambda}$, where the direct sum is over all irreps λ of G, with m_λ being the multiplicities. The projectors onto $V_\lambda^{\otimes m_\lambda} \subseteq \mathcal{H}$ can be expressed as

$$\Pi_\lambda = \frac{1}{|G|} \sum_{g\in G} \overline{\lambda(g)} U(g), \tag{D.9}$$

where $\lambda(g)$ is the character. They form an orthogonal resolution of identity. Namely,

$$\sum_{\lambda} \Pi_\lambda = \mathbb{1}, \quad \Pi_\lambda \Pi_{\lambda'} = \delta_{\lambda,\lambda'} \Pi_\lambda. \tag{D.10}$$

For a bipartition $\mathcal{H} = \mathcal{H}_A \otimes \mathcal{H}_B$, each Hilbert space $\mathcal{H}_A$ and $\mathcal{H}_B$ also decomposes into a direct sum of irrep spaces of U_A and U_B:

$$\mathcal{H}_A \simeq \bigoplus_{\lambda} V_\lambda^{\otimes m_\lambda^A}, \qquad \mathcal{H}_B \simeq \bigoplus_{\mu} V_\lambda^{\otimes m_\mu^A}, \tag{D.11}$$

where both direct sums are over all 1d irreps λ and μ of G, with multiplicities m_λ^A and m_μ^B, respectively. Note that the multiplicities can be zero. Similar constructions as (D.9) can be done for projectors onto $V_\lambda^{\otimes m_\lambda^A} \subset \mathcal{H}_A$ and $V_\lambda^{\otimes m_\lambda^B} \subset \mathcal{H}_B$.

Two irreps, λ and μ, fuse together in the tensor product to an irrep $\lambda \otimes \mu \simeq \lambda \cdot \mu$ of $U = U_A \otimes U_B$, such that

$$\mathcal{H} \simeq \left(\bigoplus_{\lambda} V_\lambda^{\otimes m_\lambda^A} \right) \otimes \left(\bigoplus_{\mu} V_\lambda^{\otimes m_\mu^A} \right) \simeq \bigoplus_{\lambda,\mu} V_{\lambda\mu}^{\otimes (m_\lambda^A + m_\mu^B)} = \bigoplus_{\Lambda} V_\Lambda^{\otimes M_\Lambda}, \tag{D.12}$$

where

$$M_\Lambda = \sum_{\lambda,\mu} \delta_{\Lambda,\lambda\mu} (m_\lambda^A + m_\mu^B) = \sum_{\lambda} (m_\lambda^A + m_{\Lambda\bar{\lambda}}^B). \tag{D.13}$$

Similarly, the projector onto the irrep Λ of $\mathcal{H}$ decomposes as,

$$\Pi_\Lambda = \sum_{\lambda} \Pi_\lambda^A \otimes \Pi_{\Lambda\bar{\lambda}}^B \tag{D.14}$$

$$= \frac{1}{|G|^2} \sum_{g,h \in G} \left(\sum_{\lambda} \overline{\lambda(g)} \lambda(h) \right) \overline{\Lambda(h)} U_A(g) \otimes U_B(h) \tag{D.15}$$

$$= \frac{1}{|G|} \sum_{g \in G} \overline{\Lambda(g)} U(g), \tag{D.16}$$

where, in the third line, we used Schur's orthogonality relations for Abelian groups:

$\sum_\lambda \overline{\lambda(g)} \lambda(h) = |G| \delta_{g,h}$. Similar identities can be derived for multipartitions.

With these tools we can prove the following lemma:

Lemma D.5. *Given a finite Abelian group G, and an homogeneous symmetry action $U = u^{\otimes N}$, with u acting on a d-dimensional on-site Hilbert space, then, as $N \to \infty$,*

$$\mathrm{Tr}[\Pi_\Lambda] = \frac{d^N}{|G/K|}(1 + O(e^{-N/\xi})), \tag{D.17}$$

where $K = \{g \in G \mid u(g) \propto \mathbb{1}\}$, and $\xi = -1/\log(\max_{g\notin K} |\mathrm{Tr}[u(g)]|/d)$.

Proof. Using that $\Pi_\Lambda = \sum_{g\in G} \overline{\Lambda(g)} U(g)$, we have

$$\mathrm{Tr}[\Pi_\Lambda] = \frac{1}{|G|}\sum_{g\in G} \overline{\Lambda(g)}\,\mathrm{Tr}[u(g)]^N. \tag{D.18}$$

From the above, we can see that the behavior of $\mathrm{Tr}[\Pi_\Lambda]$ as $N \to \infty$ is dictated by the largest values of $|\mathrm{Tr}[u(g)]|$. Since $u(g)$ is unitary, then $|\mathrm{Tr}[u(g)]| \leq d$ for all $g \in G$, and $|\mathrm{Tr}[u(g)]| = d$ if, and only if, $u(g) \propto \mathbb{1}$. The set of elements satisfying this relation form the normal subgroup $K = \{g \in G \mid u(g) \propto \mathbb{1}\}$, and the value of Λ there is fixed by the equation $\Lambda(h)\Pi_\Lambda = U(h)\Pi_\Lambda = (\mathrm{Tr}[u(h)]/d)^N \Pi_\Lambda$, where $h \in K$. In the last step, we have used that $u(h) = \mathrm{Tr}[u(h)]\mathbb{I}_d/d$, from the definition of the group K.

Finally, let $r = \max_{g\notin K} |\mathrm{Tr}[u(g)]| < d$ be the second largest absolute value of $\mathrm{Tr}[u(g)]$ in the following:

$$\mathrm{Tr}[\Pi_\Lambda] = \frac{d^N}{|G|}\sum_{h\in K} \overline{\Lambda(h)} \left(\frac{|\mathrm{Tr}[u(h)]|}{d}\right)^N + O(r^N) \tag{D.19}$$

$$= \frac{d^N}{|G/K|}(1 + O[(r/d)^N]). \tag{D.20}$$

□

If $\mathrm{Tr}[u(g)] = 0$ for all $g \in G$ satisfying $u(g) \not\propto \mathbb{1}$, then $\xi = +\infty$, and the exponentially small error above is actually zero. This is the case of $\mathbb{Z}_d$ acting on a qudit with $u(1) = Z$ or any other qudit Pauli.

Theorem D.6. *The maximally mixed symmetric state $\rho_{\beta=0,\Lambda}$ of a finite abelian group G is locally indistinguishable from the identity state, up to an exponentially small error in the trace norm:*

$$\left\|\mathrm{Tr}_{A^C}[\rho_{\beta=0,\Lambda}] - \mathbb{1}_A/d^{|A|}\right\|_1 = O(e^{-(N-|A|)/\xi}), \tag{D.21}$$

with ξ given in Lemma D.5.

Proof. Let $B = A^C$. First note that

$$\mathrm{Tr}_B[\Pi_\Lambda] = \sum_\lambda \Pi_\lambda^A \,\mathrm{Tr}_B[\Pi_{\Lambda\bar{\lambda}}^B] \tag{D.22}$$

$$= \mathbb{1}_A \frac{d^{|B|}}{|G/K|}(1 + O(e^{-|B|/\xi})), \tag{D.23}$$

where in the second line we used Lemma D.5 and $\sum_\lambda \Pi_\lambda^A = \mathbb{1}_A$.

Then, we have

$$\mathrm{Tr}_B[\rho_{\beta=0,\Lambda}] = \frac{\mathrm{Tr}_B[\Pi_\Lambda]}{\mathrm{Tr}[\Pi_\Lambda]} \tag{D.24}$$

$$= \frac{\mathbb{1}_A}{d^{|A|}}(1 + O(e^{-|B|/\xi})), \tag{D.25}$$

Hence,

$$\left\|\mathrm{Tr}_B[\rho_{\beta=0,\Lambda}] - \mathbb{1}_A/d^{|A|}\right\|_1 = O\left(\left\|\mathbb{1}_A/d^{|A|}\right\|_1 e^{-|B|/\xi}\right) = O(e^{-|B|/\xi}). \tag{D.26}$$

□

Finally, we end with a theorem that a similar result for local indistinguishability of canonical and Gibbs ensembles holds true even at finite temperatures,

Theorem D.7. *For G a finite Abelian group and $H = \sum_{i=1}^N h_i$ a local Hamiltonian with $\|h_i\| \le J$ and $|\operatorname{supp} h_i| \le k$, then*

$$\|\mathrm{Tr}_{A^C}[\rho_{\beta,\Lambda}] - \mathrm{Tr}_{A^C}[\rho_\beta]\|_1 = O(e^{-(c_\beta N - |A|)/\xi_\beta}), \tag{D.27}$$

where $\xi_\beta, c_\beta > 0$ if $\beta < \beta_*$, with ξ_β, c_β and β_* independent of system size N.

Proof. We first rewrite the difference of states above in terms of a difference in partition functions and a difference of Gibbs exponentials. For that, let $B = A^C$, and $Z_\beta := \mathrm{Tr}\left[e^{-\beta H}\right]$, $Z_{\beta,\Lambda} = \mathrm{Tr}\left[e^{-\beta H}\Pi_\Lambda\right]$:

$$\begin{aligned}\|\mathrm{Tr}_B[\rho_{\beta,\Lambda}] - \mathrm{Tr}_{A^C}[\rho_\beta]\|_1 = \Bigg\| &- \left(\frac{|G/K| Z_{\beta,\Lambda}}{Z_\beta} - 1\right) \frac{\mathrm{Tr}_B[e^{-\beta H}\Pi_\Lambda]}{Z_{\beta,\Lambda}} \\ &+ \frac{1}{Z_\beta}\Big(|G/K| \,\mathrm{Tr}_B[e^{-\beta H}\Pi_\Lambda] - \mathrm{Tr}_B[e^{-\beta H}]\Big)\Bigg\|_1 \\ \le\; & Z_\beta^{-1}\big||G/K| Z_{\beta,\Lambda} - Z_\beta\big| \,\|\mathrm{Tr}_B[\rho_{\beta,\Lambda}]\|_1 \\ &+ Z_\beta^{-1}\big\||G/K| \,\mathrm{Tr}_B[e^{-\beta H}\Pi_\Lambda] - \mathrm{Tr}_B[e^{-\beta H}]\big\|_1.\end{aligned} \tag{D.28}$$

The first term above is nothing but the second term specialized to $A = \emptyset$. Hence, we will focus on bounding the latter. Without loss of generality, we can assume $H \leq 0$ by adding to it a constant term proportional to the identity operator that will not change ρ_β nor $\rho_{\beta,\Lambda}$. In that case, $Z_\beta = \mathrm{Tr}\big[e^{-\beta H}\big] \geq d^N$. By applying a high-temperature expansion of the exponentials, we arrive at

$$Z_\beta^{-1} \| |G/K| \, \mathrm{Tr}_B[e^{-\beta H}\Pi_\Lambda] - \mathrm{Tr}_B[e^{-\beta H}] \|_1 \leq d^{-N} \sum_{n=0}^{\infty} \frac{\beta^n}{n!} \| |G/K| \, \mathrm{Tr}_B[H^n \Pi_\Lambda] - \mathrm{Tr}_B[H^n] \|_1. \tag{D.29}$$

We now focus on each term of the sum above,

$$\begin{aligned} &\| |G/K| \, \mathrm{Tr}_B[H^n \Pi_\Lambda] - \mathrm{Tr}_B[H^n] \|_1 \\ &\qquad \leq \sum_{i_1,\ldots,i_n} \left\| |G/K| \, \mathrm{Tr}_B \left[\prod_{\ell=1}^{n} h_{i_\ell} \Pi_\Lambda \right] - \mathrm{Tr}_B \left[\prod_{\ell=1}^{n} h_{i_\ell} \right] \right\|_1 \end{aligned} \tag{D.30}$$

At this stage, it is useful to introduce a notation for the region with non-trivial action of $\prod_{\ell=1}^n h_{i_\ell}$, $S_{\{i_\ell\}} = \cup_{\ell=1}^n \operatorname{supp} h_{i_\ell}$. We decompose the trace over B into trace over $B \cap S_{\{i_\ell\}}$ and $B \cap S_{\{i_\ell\}}^C$. We also decompose $\Pi_\Lambda = \sum_\lambda \Pi_\lambda^R \Pi_{\Lambda\lambda}^{R^C}$ for $R = A \cup S_{\{i_\ell\}}$, such that $R^C \equiv B \cap S_{\{i_\ell\}}^C$.

We also use the fact that $\sum_\lambda \Pi_\lambda^R = \mathbb{1}_R$ to simplify the second term in Eq. (D.30):

$$
\begin{aligned}
&\| |G/K| \operatorname{Tr}_B[H^n \Pi_\Lambda] - \operatorname{Tr}_B[H^n] \|_1 \\
&\leq \sum_{i_1,\ldots,i_n} \left\| \sum_\lambda \operatorname{Tr}_{B\cap S_{\{i_\ell\}}} \left[\prod_{\ell=1}^n h_{i_\ell} \Pi_\lambda^{A\cup S_{\{i_\ell\}}} \right] \left(|G/K| \operatorname{Tr}_{B\cap S^C_{\{i_\ell\}}} \left[\Pi_{\Lambda\bar{\lambda}}^{B\cap S^C_{\{i_\ell\}}} \right] - \operatorname{Tr}_{B\cap S^C_{\{i_\ell\}}} \left[\mathbb{I}_{B\cap S^C_{\{i_\ell\}}} \right] \right) \right\|_1 && \text{(D.31)} \\
&= \sum_{i_1,\ldots,i_n} \sum_\lambda \left\| \operatorname{Tr}_{B\cap S_{\{i_\ell\}}} \left[\prod_{\ell=1}^n h_{i_\ell} \Pi_\lambda^{A\cup S_{\{i_\ell\}}} \right] \right\|_1 \left| |G/K| \operatorname{Tr}_{B\cap S^C_{\{i_\ell\}}} \left[\Pi_{\Lambda\bar{\lambda}}^{B\cap S^C_{\{i_\ell\}}} \right] - d^{|B\cap S^C_{\{i_\ell\}}|} \right| && \text{(D.32)} \\
&\leq C \sum_{i_1,\ldots,i_n} \sum_\lambda \left\| \operatorname{Tr}_{B\cap S_{\{i_\ell\}}} \left[\prod_{\ell=1}^n h_{i_\ell} \Pi_\lambda^{A\cup S_{\{i_\ell\}}} \right] \right\|_1 r^{|B\cap S^C_{\{i_\ell\}}|} \text{ (using Lemma D.5)} && \text{(D.33)} \\
&\leq C \sum_{i_1,\ldots,i_n} \sum_\lambda \left\| \prod_{\ell=1}^n h_{i_\ell} \Pi_\lambda^{A\cup S_{\{i_\ell\}}} \right\|_1 r^{|B|-\max(|B|,kn)} && \text{(D.34)} \\
&= C \sum_{i_1,\ldots,i_n} \left\| \prod_{\ell=1}^n h_{i_\ell} \right\|_1 r^{|B|-\max(|B|,kn)} && \text{(D.35)} \\
&\leq C \sum_{i_1,\ldots,i_n} d^{|A|+\max(|B|,kn)} J^n r^{|B|-\max(|B|,kn)} && \text{(D.36)} \\
&\leq C \begin{cases} r^N (d/r)^{|A|} (JNd^k/r^k)^n, & \text{if } kn < |B|, \\ d^N (JN)^n, & \text{if } kn \geq |B|. \end{cases} && \text{(D.37)}
\end{aligned}
$$

where in the third line we have used Lemma D.5 in the form $||G/K| \operatorname{Tr}\left[\Pi_\lambda^R\right] - d^{|R|}| < C r^{|R|}$, with C a constant and, crucially, $r < d$; and in the fifth line we used that $\prod_{\ell=1}^n h_{i_\ell}$ is block diagonal with respect to the subspaces projected by $\Pi_\lambda^{A\cup S_{\{i_\ell\}}}$ because the product commutes

with the symmetry acting on $A \cup S_{\{i_\ell\}}$. Putting it back in the original sum (Eq. (D.29)),

$$Z_\beta^{-1} \| |G/K| \operatorname{Tr}_B[e^{-\beta H}\Pi_\Lambda] - \operatorname{Tr}_B[e^{-\beta H}] \|_1 \tag{D.38}$$

$$\leq C \left((r/d)^{N-|A|} \sum_{n=0}^{\lceil |B|/k \rceil - 1} \frac{(\beta J N d^k/r^k)^n}{n!} + \sum_{n=\lceil |B|/k \rceil}^{\infty} \frac{(\beta J N)^n}{n!} \right) \tag{D.39}$$

$$\leq C \left((r/d)^{N-|A|} \sum_{n=0}^{\infty} \frac{(\beta J N d^k/r^k)^n}{n!} + e^{\beta J N} \frac{(\beta J N)^{\lceil |B|/k \rceil}}{\lceil |B|/k \rceil !} \right) \tag{D.40}$$

$$\leq C \left((r/d)^{N-|A|} e^{\beta J N d^k/r^k} + e^{\beta J N} \left(\frac{e \beta J N}{\lceil |B|/k \rceil} \right)^{\lceil |B|/k \rceil} \right) \tag{D.41}$$

$$\leq C (e^{-(\ln(d/r) - \beta J d^k/r^k) N + \ln(d/r)|A|} + e^{(\beta J + \ln(e\beta J k))N + (-\ln(e\beta J k)+1)|A|}) \tag{D.42}$$

$$\leq C (e^{-(\ln(d/r) - \beta J d^k/r^k) N + \ln(d/r)|A|} + e^{\beta J N} (e\beta J k)^{N-|A|} (1 - |A|/N)^{-(N-|A|)}) \tag{D.43}$$

$$= O(e^{-(c_\beta N - |A|)/\xi_\beta}), \tag{D.44}$$

where the values of ξ_β and c_β depend if $r = 0$ or not: if $r = 0$, then the first term above vanishes, and so $\xi_\beta^{-1} = 1 - \ln(e\beta J k)$ and $c_\beta = -\xi_\beta(\beta J + \ln(e\beta J k))$; and if $r > 0$, then, at small enough β, the first term dominates the second, and we have $\xi_\beta^{-1} = \ln(d/r)$ and $c_\beta = 1 - \xi_\beta \beta J d^k/r^k$. Importantly, in both cases we have $\xi_\beta, c_\beta > 0$ if $\beta < \beta_*$ for β_* small enough.

Since this bound is tighter the smaller region A is, then the first term of Eq. (D.28) is smaller than the second one, and we arrive at the conclusion. □

D.5 Equivalence of SEC and EC for finite Abelian groups

Here, we find a sufficient condition on the symmetry multiplicities that allows us to establish equivalence between the entangling condition (EC) and (SEC) for finite Abelian groups. First, a technical lemma.

Lemma D.8. *Consider a Hamiltonian defined on a Hilbert space with the tripartition $A|B|C$, $\mathcal{H} = \mathcal{H}_A \otimes \mathcal{H}_B \otimes \mathcal{H}_C$, $H = H_{AB} \otimes \mathbb{1}_C + \mathbb{1}_A \otimes H_{BC}$, such that B acts as a buffer region between A and C with no global interaction term nor any local interaction term between A and C. Furthermore, each term H_{AB} and H_{BC} is assumed to be individually symmetric: $\forall g \in G, [U(g), H_{AB}] = [U(g), H_{BC}] = 0$. If all irreps λ that appear in AB*

(i.e. $m_\lambda^{AB} \neq 0$) contribute to the total irrep Λ (i.e. $m_{\Lambda\bar{\lambda}}^{C} \neq 0$), then we have an equivalence between the EC and SEC,

$$[U_A(g), H]\Pi_\Lambda \neq 0 \iff [U_A(g), H] \neq 0. \tag{D.45}$$

Proof. The ($\Rightarrow$) direction is trivial. For the ($\Leftarrow$) direction, we prove their negations are equivalent. We first right-multiply the RHS $[U_A(g), H]\Pi_\Lambda = 0$ by Π_λ^{AB} where λ is any irrep that appears in $\mathcal{H}_{AB}$, i.e. $m_\lambda^{AB} \neq 0$:

$$0 = [U_A(g), H]\Pi_\Lambda\Pi_\lambda^{AB} = [U_A(g), H]\Pi_\lambda^{AB}\Pi_{\Lambda\bar{\lambda}}^{C} = [U_A(g), H_{AB}]\Pi_\lambda^{AB}\Pi_{\Lambda\bar{\lambda}}^{C}. \tag{D.46}$$

Then, we trace out the C subsystem:

$$0 = [U_A(g), H_{AB}]\Pi_\lambda^{AB} \otimes \mathrm{Tr}\big[\Pi_{\Lambda\bar{\lambda}}^{C}\big] = m_{\Lambda\bar{\lambda}}^{C}[U_A(g), H_{AB}]\Pi_\lambda^{AB}. \tag{D.47}$$

By assumption, $m_\lambda^{AB} \neq 0$ implies $m_{\Lambda\bar{\lambda}}^{C} \neq 0$, which allows us to get rid of $m_{\Lambda\bar{\lambda}}^{C}$ from the equation above. Summing over λ, we finally arrive at

$$\begin{aligned} 0 &= \sum_{\lambda, m_\lambda^{AB} \neq 0} [U_A(g), H_{AB}]\Pi_\lambda^{AB} && \text{(D.48)} \\ &= \sum_{\lambda} \Pi_\lambda^{AB}[U_A(g), H_{AB}]\Pi_\lambda^{AB} && \text{(D.49)} \\ &= \sum_{\lambda,\lambda'} \Pi_\lambda^{AB}[U_A(g), H_{AB}]\Pi_{\lambda'}^{AB} && \text{(D.50)} \\ &= [U_A(g), H_{AB}], && \text{(D.51)} \end{aligned}$$

which, after restoring the C system by tensoring with $\mathbb{1}_C$, we have $[U_A(g), H] = 0$. □

In a multipartite system $\mathcal{H} = \bigotimes_{i\in M} \mathcal{H}_i$ with a Hamiltonian $H = \sum_{X\subseteq M} h_X$ that is a sum of local terms h_X, $\mathrm{supp}(h_X) = X$, we usually require them to be individually symmetric, $\forall X \subseteq M, \forall g \in G, [U(g), h_X] = [U_X(g), h_X] = 0$. Here, locality means that for all regions A with size $|A| \leq N/2$, there exists a buffer region B with size $|B| < N/2$ (normally, $|B| \ll N$) for which H has no direct interaction between A and $C = (AB)^C$, i.e. $X \cap A \cap C \neq 0 \Rightarrow h_X = 0$. Then the Lemma above can be applied for such A and B regions if its conditions are met. In such case, we know that $\forall A \subseteq M, |A| < N/2 \Rightarrow \exists g \in G, [U_A(g), H] \neq 0$ is equivalent to $\rho_{\beta,\Lambda}$ having nonzero negativity for all bipartitions at high enough temperatures.

For finite Abelian symmetry groups G acting faithfully, then every irrep of G should appear for large enough regions C. Moreover, "large enough" here should depend only on

the group itself, and not on the representation or system size. For such cases, then the above lemma is useful to prove POE purposes for any finite-range Hamiltonian. Let us make this intuition more precise, starting with the following lemma:

Lemma D.9 (Stable power of generating set)**.** *Let H be a finite group (not necessarily Abelian) and $S \subseteq H$ a generating set of $H = \langle S \rangle$ containing the identity element $e \in S$. With S^n being the set of all products of n elements of S, we have*

$$\{e\} = S^0 \subseteq S^1 \subseteq S^2 \subseteq \cdots \subseteq S^{|H|-1} = S^{|H|} = \cdots = H. \tag{D.52}$$

Proof. First, note that $m \leq n \implies S^m \subseteq S^n$ since $S^n = S^m \cdot S^{n-m} \supseteq S^m \cdot \{e\} = S^m$. Moreover, since H is finite and S generates it, then, for large enough n, $S^n = H$. Hence, it only remains to prove that $\forall n \geq |H|, S^{n-1} = S^n$.

Suppose, by contradiction, that there is an element $h \in S^n \setminus S^{n-1}$. Since $h \in S^n$, there are $s_1, \ldots, s_n \in S$ such that $h = s_1 s_2 \cdots s_n$. Now, recursively define the partial product sequence $(h_k)_{k=0}^n$ by $h_0 = e$ and $h_k = h_{k-1} s_k$. Crucially, such sequence satisfies $h_k \in S^k$ and $h_k^{-1} h \in S^{n-k}$.

We will prove that the assumption $h \notin S^{n-1}$ implies all elements h_k are distinct from each other. Otherwise, if $h_k = h_l$ for $0 \leq k < l \leq n$, then

$$h = h_l(h_l^{-1} h) = h_k(h_l^{-1} h) \in S^k \cdot S^{n-l} \subseteq S^{n-1}, \tag{D.53}$$

contradicting the assumption $h \notin S^{n-1}$. However, H has only $|H|$ elements, so the $n + 1 \geq |H| + 1$ elements of $(h_k)_{k=0}^n$ cannot be all distinct. Thus, the original hypothesis of $\exists h \in H, h \in S^n \setminus S^{n-1}$ is incorrect, and we have $S^{n-1} = S^n$. □

The condition $e \in S$ is important. For a counterexample, consider $H = \mathbb{Z}_2 \times \mathbb{Z}_2$ and $S = \{(1,0), (1,1)\}$. Even though S generates H, its powers alternate between $S = S^3 = \cdots = \{(1,0), (1,1)\}$ and $S^2 = S^4 = \cdots = \{(0,0), (0,1)\}$, never stabilizing to the full group H.

We can apply the lemma above to a subgroup of $\widehat{G}$, the group of 1d irreps of G (also called the Pontryagin dual of G), generated by the set of irreps appearing in the on-site representation u, here assumed to be site-independent. Let us frame it in the context of spin chains for simplicity, but a similar result will be valid for any finite-range Hamiltonian in any bounded-degree graph (which includes higher dimensional lattices).

Theorem D.10. *Let G be a finite Abelian group. Consider a spin chain with periodic boundary conditions and homogeneous symmetry action $U = u^{\otimes N} : G \to U(\mathcal{H})$, with*

u containing the trivial irrep in its decomposition. Further assume that the terms h_i of the Hamiltonian $H = \sum_i h_i$ are supported around site $i \ni \mathrm{supp}(h_i)$, have finite range d, $\mathrm{diam}(\mathrm{supp}(h_i)) \leq d$, and are individually symmetric under U. For any global irrep Λ and any connected region A (an interval) separating $H = H_A \otimes \mathbb{1}_{A^C} + \mathbb{1}_A \otimes H_{A^C} + H_{AA^C}$ with $|A| \leq N - (2d + |G| - 3) = N - O(1)$, we have

$$[U_A, H]\Pi_\Lambda \neq 0 \iff [U_A, H] \neq 0. \tag{D.54}$$

Proof. Let $S = \{\lambda \in \widehat{G} \mid \lambda \in u\}$, where "$\lambda \in u$" means that λ appears in the irrep decomposition of the on-site representation u with nonzero multiplicity. By assumption $e \in S$. Then, it is easy to see that $S^n = \{\lambda \in \widehat{G} \mid \lambda \in u^{\otimes n}\}$. Lemma D.9 guarantees that all irreps that will ever appear in any tensor product of u form the subgroup $\langle S\rangle \leq \widehat{G}$, and that all regions C with size $|C| \geq |\widehat{G}| - 1 = |G| - 1$ will contain them, i.e. $S^{|C|} = \langle S\rangle$. Here, we used that if G is a finite Abelian group, then $G \simeq \widehat{G}$. That follows from $\widehat{\mathbb{Z}_n} \simeq \mathbb{Z}_n$ and the fact that G is a direct product of cyclic groups. For such regions C, we can apply Lemma D.8 if B is large enough to separate the interval A from C. That is because both the irrep λ of AB with multiplicity $m_\lambda^{AB} \neq 0$ is an element of $S^{|AB|} \subseteq \langle S\rangle$, and the total irrep Λ is an element of $S^N = \langle S\rangle$, which means that $\Lambda\overline{\lambda} \in \langle S\rangle = S^{|C|}$ appears in C. For that, we choose $B = \{i \in A^C \mid \mathrm{dist}(A, i) \leq d - 1\}$, and $C = (AB)^C$. Given that A is connected, these three regions separate Hamiltonian into $H = \mathbb{1}_A \otimes H_{BC} + H_{AB} \otimes \mathbb{1}_C$, thus satisfy all of the assumptions of Lemma D.8. Finally, the condition $|C| \geq |G| - 1$ translates to $|A| = N - |B| - |C| \leq N - 2(d - 1) - (|G| - 1)$. □

As an example, let us consider a nearest-neighbor Hamiltonian $H = \sum_i h_{i,i+1}$ with homogeneous on-site $\mathbb{Z}_2$ symmetry given by $\prod_i u_i$, $u_i^2 = \mathbb{1}_i$ (so $d = 2$ and $|G| = 2$). If the terms $h_{i,i+1}$ are $\mathbb{Z}_2$ symmetric, but not diagonal in the u_i basis, then they will not commute with the partial action of the symmetry, i.e. $[u_i, h_{i,i+1}] \neq 0$. Under these assumptions, the corollary above implies that any canonical ensemble state will not have SDOE for all bipartitions into two intervals, if the total system size N is larger than 5. That is because $\lfloor N/2 \rfloor \leq N - (2d + |G| - 3) = N - 3$ for $N \geq 5$.

Note that the finiteness of G is important for the above. For an example of an infinite Abelian symmetry with tighter constraints on the regions A for which Lemma D.8 can be applied, consider an homogeneous $G = U(1)$ symmetry action with on-site charges 0 and +1, which is the case for $U(1)$ acting on qubits as $\theta \mapsto \mathrm{diag}(1, e^{i\theta})$. Then, the condition of Lemma D.8 is equivalent to the following condition on the total charge n ($\Lambda(\theta) = \exp(in\theta)$): $|AB| \leq n \leq N - |AB|$. In particular, the Lemma can be applied to any finite-range Hamiltonian constrained to a total charge $n \approx N - O(1)$, if only entanglement

of finite regions $|A| = O(1)$ is relevant. For larger regions, such as $|A| = N/2 - O(1)$, the effect of the global charge constraint can be noticeable, which does not happen for finite Abelian groups (if N is large enough).

D.6 Classification of irreps

Here we provide a detailed proof of Theorem 5.4 presented in the main text. The definition of semiuniform irreps and the theorem is quoted below for the ease of readers.

Definition D.11. A global irrep Λ of a N-partite system is **uniform** if its projector Π_λ is a uniform tensor product of an on-site irrep λ: $\Pi_\Lambda = \bigotimes_{i=1}^N \Pi_\lambda^{(i)}$. Λ is **semiuniform** if its projector Π_Λ is a direct sum of uniform tensor products: $\Pi_\Lambda = \sum_\alpha \bigotimes_{i=1}^N \Pi_{\lambda_\alpha}^{(i)}$.

Theorem D.12. *Consider a multipartite system with homogeneous symmetry action U. If Λ*

1. *is not semiuniform, then for any pair of sites (i,j), there exists a two-body symmetric Hamiltonian $v_{i,j}$ supported on $\{i,j\}$ that entangles any bipartition that separates site i from site j. Moreover, given a connected interaction graph, the uniform nearest-neighbor perturbation $\sum_{\langle i,j\rangle} v_{i,j}$ entangling all bipartitions $A|B$, with $A, B \neq \emptyset$.*
2. *is semiuniform, then there is no entangling Hamiltonian with support over $N-1$ sites or fewer, and for all $(N-1)$-local Hamiltonian, $\rho_{\beta,\Lambda}$ is fully separable when ρ_β is fully separable.*

Proof. (A) Since Λ is not semiuniform, then there is a tensor product projector term $\bigotimes_{i=1}^N \Pi_{\lambda_i}^{(i)}$ in the decomposition of Π_Λ with at least two distinct irreps over sites k and l: $\lambda \neq \lambda'$. Hence,

$$\Pi_\lambda^{(l)} \Pi_{\lambda'}^{(k)} \Pi_\Lambda \neq 0.$$

Since Π_Λ has weak permutation symmetry from the homogeneity of the symmetry action, then we can permute sites (k,l) with (i,j) and (j,i) on the equation above, thus proving that $\Pi_\lambda^{(i)} \Pi_{\lambda'}^{(j)} \Pi_\Lambda, \Pi_{\lambda'}^{(i)} \Pi_\lambda^{(j)} \Pi_\Lambda \neq 0$.

Thus there exist states $|\lambda\rangle_i$ and $|\lambda'\rangle_i$ of $\mathcal{H}_i$ that transform as irreps λ and λ' under the symmetry, respectively, and similarly with two states $|\lambda\rangle_j$ and $|\lambda'\rangle_j$ of $\mathcal{H}_j$. Then, let $\sigma_i^+ := {}_i|\lambda'\rangle\langle\lambda|_i \in \mathcal{L}(\mathcal{H}_i)$ and $\sigma_j^+ := {}_j|\lambda'\rangle\langle\lambda|_j \in \mathcal{L}(\mathcal{H}_j)$. Denoting their Hermitian conjugates by $\sigma_i^- := (\sigma_i^+)^\dagger$ and $\sigma_j^- := (\sigma_j^+)^\dagger$, we choose $v_{i,j} := \sigma_i^+ \sigma_j^- + \sigma_i^- \sigma_j^+$, which is an entangling Hamiltonian because:

- $v_{i,j}$ is symmetric under $U = U_i \otimes U_j \otimes_{k \neq i,j} U_k$:

$$U(g) v_{i,j} U(g)^\dagger = \lambda'(g)\lambda(g)\sigma_i^+ \sigma_j^- \overline{\lambda(g)\lambda'(g)} + \text{h.c.} = v_{i,j}, \tag{D.55}$$

- $v_{i,j}$ satisfies the entangling condition EC of Theorem 5.2. To see this, note that the assumptions imply the existence of two states, $|\Psi\rangle = |\lambda\rangle_i \otimes |\lambda'\rangle_j \otimes |\psi\rangle_{\setminus i,j}$ and $|\Psi'\rangle = |\lambda'\rangle_i \otimes |\lambda\rangle_j \otimes |\psi\rangle_{\setminus i,j}$ that transform as the global irrep Λ. Moreover, since $\lambda \neq \lambda'$, there exists $g \in G$ such that $\lambda(g) \neq \lambda'(g)$. Now we establish a nonzero matrix element of $[U_i(g), H]\Pi_\Lambda$:

$$\langle \Psi' | [U_i(g), v_{i,j}] \Pi_\Lambda | \Psi \rangle = (\lambda'(g) - \lambda(g)) \langle \Psi' | v_{i,j} | \Psi \rangle = \lambda'(g) - \lambda(g) \neq 0. \tag{D.56}$$

We now turn to the uniform nearest-neighbors perturbation $V = \sum_{\{i,j\} \in E} v_{i,j}$, where E is the set of edges of a connected graph $\mathcal{G} = (V, E)$. By contradiction, let us assume it is not entangling for a certain region $A \neq \emptyset, V$.

In what follows, let $\partial A := \{\{i,j\} \in E \mid i \in A, j \notin A\}$, and $g \in G$ be a group element for which $\lambda'(g) \neq \lambda(g)$. Since V is not entangling, then we must have,

$$0 = [U_A(g), V]\Pi_\Lambda = (\lambda'(g) - \lambda(g))\Pi_\Lambda \sum_{\{i,j\} \in \partial A} \sigma_i^+ \sigma_j^- + \text{h.c.}. \tag{D.57}$$

Since $A \neq \emptyset, V$ and $\mathcal{G}$ is connected, the sum above contains at least one term corresponding to an edge $\{i^*, j^*\} \in \partial A$. Now, let $|\Psi\rangle = |\lambda\rangle_{i^*} \otimes |\lambda'\rangle_{j^*} \bigotimes_{l \neq i^*, j^*} |\psi_l\rangle_l$ and $|\Psi'\rangle = |\lambda'\rangle_{i^*} \otimes |\lambda\rangle_{j^*} \bigotimes_{l \neq i^*, j^*} |\psi_l\rangle_l$ be tensor product states contained within Λ, i.e. $\Pi_\Lambda |\Psi\rangle = |\Psi\rangle$. Then,

$$\begin{aligned} 0 &= \langle \Psi' | (\lambda'(g) - \lambda(g))\Pi_\Lambda \sum_{(i,j) \in \partial A} \sigma_i^+ \sigma_j^- + \text{h.c.} | \Psi \rangle & \text{(D.58)} \\ &= (\lambda'(g) - \lambda(g)) \langle \Psi' | \sigma_{i^*}^+ \sigma_{j^*}^- | \Psi \rangle & \text{(D.59)} \\ &= \lambda'(g) - \lambda(g), & \text{(D.60)} \end{aligned}$$

a contradiction. Here, only the term $\sigma_{i^*}^+ \sigma_{j^*}^-$ survives because the others are either $\sigma_{j^*}^+ \sigma_{i^*}^-$ or they contain at least one vertex, say $l \in V$, that is not i^* or j^*, and thus will get a diagonal contribution corresponding to the state $|\psi_l\rangle_l$. In both cases, the matrix element is zero by construction.

Thus, $V = \sum_{\{i,j\} \in E} v_{i,j}$ is an entangling Hamiltonian for all regions $A \neq \emptyset, V$, as desired.

(B) Say Λ is semiuniform with projector $\Pi_\Lambda = \sum_\alpha \Pi_\alpha = \sum_\alpha \bigotimes_{i=1}^N \Pi_{\lambda_\alpha}^{(i)}$, where $\Pi_\alpha = \bigotimes_{i=1}^N \Pi_{\lambda_\alpha}^{(i)}$ are orthogonal tensor product projectors. Let H be an arbitrary Hamiltonian that is G-symmetric and supported over the $N-1$ first sites, without loss of generality. Then,

$$H\Pi_\Lambda = \Pi_\Lambda H \Pi_\Lambda \tag{D.61}$$

$$= \sum_{\alpha,\beta} \left[\left(\bigotimes_{i=1}^{N-1} \Pi_{\lambda_\alpha}^{(i)} \right) H \left(\bigotimes_{i=1}^{N-1} \Pi_{\lambda_\beta}^{(i)} \right) \right] \otimes \Pi_{\lambda_\alpha}^N \Pi_{\lambda_\beta}^N \tag{D.62}$$

$$= \sum_{\alpha} \left(\bigotimes_{i=1}^{N} \Pi_{\lambda_\alpha}^{(i)} \right) H \left(\bigotimes_{i=1}^{N} \Pi_{\lambda_\alpha}^{(i)} \right). \tag{D.63}$$

From Eq. (D.63), we have that $\alpha \neq \beta \Rightarrow \Pi_\alpha (H\Pi_\Lambda)\Pi_\beta = 0$, which also implies $\Pi_\alpha e^{-\beta' H \Pi_\Lambda} \Pi_\beta = 0$. As such, if $\rho_{\beta'} \propto e^{-\beta' H}$ is separable with decomposition $\rho_{\beta'} = \sum_m p_m |\Psi_m\rangle\langle\Psi_m|$, with $|\Psi_m\rangle = \bigotimes_{i=1}^N |\psi_m^{(i)}\rangle$ product state, then

$$\rho_{\beta',\Lambda} \propto e^{-\beta' H} \Pi_\Lambda \tag{D.64}$$

$$= \Pi_\Lambda e^{-\beta' H \Pi_\Lambda} \Pi_\Lambda \tag{D.65}$$

$$= \sum_{\alpha,\beta} \Pi_\alpha e^{-\beta' H \Pi_\Lambda} \Pi_\beta \tag{D.66}$$

$$= \sum_{\alpha} \Pi_\alpha e^{-\beta' H} \Pi_\alpha \tag{D.67}$$

$$\propto \sum_{\alpha,m} p_m (\Pi_\alpha |\Psi_m\rangle)(\langle\Psi_m| \Pi_\alpha), \tag{D.68}$$

which gives $\rho_{\beta,\Lambda}$ a fully separable decomposition into the product states

$$\Pi_\alpha |\Psi_m\rangle = \bigotimes_{i=1}^N (\Pi_{\lambda_\alpha}^{(i)} |\psi_m^{(i)}\rangle)$$

.

□

Now, it remains to prove that semiuniform irreps are rare among the set of all global irreps:

Theorem D.13. *Consider a N-partite system with an Abelian symmetry group G that 1) does not depend on N, 2) acts homogeneously (i.e. $U = u^{\otimes N}$), and 3) acts nontrivially on each site (i.e. $\exists g \in G, u(g) \not\propto \mathbb{1}$, or equivalently there are at least two distinct on-site irreps). Then, the fraction of semiuniform global irreps goes to zero in the thermodynamic limit as $O(2^d/N)$, where d is the on-site Hilbert space dimension. Furthermore, if G is finite, there are no semiuniform global irreps for $N > |G|$.*

Proof. We separate the analysis in two cases: either $\lambda\overline{\mu} \in \widehat{G}$ has finite order for all on-site irreps $\lambda, \mu \in u$, or there is a pair $(\lambda_\infty, \mu_\infty)$ with $\lambda_\infty\overline{\mu_\infty}$ having infinite order.

1. In the first case, define $n < \infty$ be the least common multiple of the orders of all such $\lambda\overline{\mu}$. Note that this is always the case for *finite* symmetry groups G, with $n \leq |\widehat{G}| = |G|$. We will prove that there are no semiuniform global irreps for $N > n$, thus trivially satisfying the fraction bound.

   Indeed, if Λ is semiuniform with $\Pi_\Lambda \neq 0$, there is at least one λ^* one of the irreps that appear as a nonzero uniform term $\bigotimes_{i=1}^N \Pi_{\lambda^*}^{(i)}$ in the sum. Furthermore, since u acts nontrivially, let $\mu \neq \lambda^*$ be another on-site irrep, distinct from λ^*, with $\Pi_\mu^{(i)} \neq 0$. Then, we can construct a *non-uniform* irrep vector $\vec{\lambda}^{NU}$ given by $\lambda_i^{NU} = \mu$ if $1 \leq i \leq n$ and $\lambda_i^{NU} = \lambda^*$ if $n < i \leq N$. This irrep appears in the decomposition of Π_Λ because $\mu^n(\lambda^*)^{N-n} = (\lambda^*)^N = \Lambda$, where we have used that $(\lambda^*)^n = \mu^n = 1$, the trivial irrep. This, however, contradicts the assumption of semiuniformity of Λ.

2. In the latter case — of $\lambda_\infty\overline{\mu_\infty}$ having infinite order — we can lower bound the number of global irreps by choosing some of the on-site irreps to be λ_∞, and some to be μ_∞. As a result, we form a set $\{\lambda_\infty^m \mu_\infty^{N-m}\}_{m=0}^N$ of $N+1$ global irreps that are *pairwise distinct* precisely because $\lambda_\infty\overline{\mu_\infty}$ has infinite order.

   To upper bound the number of *semiuniform* global irreps, first note that there are at most d distinct on-site irreps contained in the decomposition of u. Since a semiuniform irrep is a sum over a subset of on-site irreps, there are at most 2^d of them. Hence, the fraction of semiuniform global irreps is bounded by $2^d/(N+1) < 2^d/N$.

□

For some infinite groups, including $U(1)$, the fraction of irreps can be further constrained to be at most $d/(N+1)$ due to the following result:

Proposition D.14. *If $\widehat{G}$ is torsion-less (e.g. $G = U(1)$), then all semiuniform global irreps Λ are uniform.*

Proof. If there are two uniform tensor products, $\bigotimes_i \Pi_\lambda^{(i)}$ and $\bigotimes_i \Pi_\mu^{(i)}$, contributing to Π_Λ, then this means $(\lambda\overline{\mu})^N = \Lambda\overline{\Lambda} = 1$. Since $\widehat{G}$ is torsion-less, then $\lambda = \mu$. By repeating this argument, we conclude that Λ is uniform. □

D.7 Equivalence of EC and NC

Here we prove that the conditions EC (eq. 5.2) and NC (eq. 5.3) are equivalent. This shows that whenever a symmetric Hamiltonian and irrep satisfies EC to lead to entangled canonical ensembles, the latter must also be negative partial transpose, i.e., their entanglement negativity must be non-zero. We prove the equivalence through the following theorem:

Theorem D.15. *The following are equivalent:*

1. $(\mathbb{1} - \Pi_\Lambda)(H\Pi_\Lambda)^{T_A}(\mathbb{1} - \Pi_\Lambda) = 0$.
2. $H\Pi_\Lambda = \sum_\lambda \Pi_\lambda^A (H\Pi_\Lambda)\Pi_\lambda^A = \sum_\lambda \Pi_\lambda^A \Pi_{\Lambda\overline{\lambda}}^B H \Pi_\lambda^A \Pi_{\Lambda\overline{\lambda}}^B$.
3. *For all* $g \in G$, $[U_A(g), H\Pi_\Lambda] = [U_B(g), H\Pi_\Lambda] = 0$.

Proof.

- $(i) \Rightarrow (ii)$ First, we have

$$(H\Pi_\Lambda)^{T_A} = \sum_\lambda \Pi_\lambda^A H^{T_A} \Pi_{\Lambda\overline{\lambda}}^B, \tag{D.69}$$

where we have used $(\Pi_\lambda^A)^T = \Pi_\lambda^A$, which is true as we choose to do the partial transpose for A and B with respect to a basis that diagonalizes U_A and U_B, such that $\forall g \in G, U_A(g)^{T_A} = U_A(g)$ and $U_B(g)^{T_B} = U_B(g)$. Then,

$$0 = [(\mathbb{1} - \Pi_\Lambda)(H\Pi_\Lambda)^{T_A}(\mathbb{1} - \Pi_\Lambda)]^{T_A} \tag{D.70}$$

$$\left[\sum_\lambda \Pi_\lambda^A H^{T_A} \Pi_{\Lambda\overline{\lambda}}^B + \sum_\lambda \Pi_{\Lambda\overline{\lambda}}^B H^{T_A} \Pi_\lambda^A - \sum_\lambda \Pi_{\Lambda\overline{\lambda}}^B H^{T_A} \Pi_{\Lambda\overline{\lambda}}^B - \sum_\lambda \Pi_\lambda^A H^{T_A} \Pi_\lambda^A\right]^{T_A} \tag{D.71}$$

$$= H\Pi_\Lambda + \Pi_\Lambda H - \sum_\lambda \Pi_{\Lambda\overline{\lambda}}^B H \Pi_{\Lambda\overline{\lambda}}^B - \sum_\lambda \Pi_\lambda^A H \Pi_\lambda^A. \tag{D.72}$$

By multiplying the equation above with Π_Λ and using that $[H, \Pi_\Lambda] = 0$, we have

$$H\Pi_\Lambda = \frac{1}{2}\left(\sum_\lambda \Pi_\Lambda \Pi_{\Lambda\overline{\lambda}}^B H \Pi_{\Lambda\overline{\lambda}}^B \Pi_\Lambda + \sum_\lambda \Pi_\lambda^A H \Pi_\Lambda \Pi_\lambda^A\right) = \sum_\lambda \Pi_\lambda^A (H\Pi_\Lambda)\Pi_\lambda^A. \tag{D.73}$$

- $(ii) \Rightarrow (iii)$ From assertion (ii), we know that $H\Pi_\Lambda$ is block diagonal with respect to the irreps of U_A. Thus, by the converse of Schur's lemma, $H\Pi_\Lambda$ commutes with all symmetry operators $U_A(g)$, for all $g \in G$. More explicitly, from $U_A(g)\Pi^A_\lambda = \Pi^A_\lambda U_A(g) = \lambda(g)\Pi^A_\lambda$, we have

$$[U_A(g), H\Pi_\Lambda] = \sum_\lambda \lambda(g)\Pi^A_\lambda(H\Pi_\Lambda)\Pi^A_\lambda - \sum_\lambda \Pi^A_\lambda(H\Pi_\Lambda)\Pi^A_\lambda\lambda(g) = 0. \tag{D.74}$$

- $(iii) \Rightarrow (i)$ Condition (iii) also implies that for all irreps λ and μ, $[\Pi^A_\lambda, H\Pi_\Lambda] = [\Pi^B_\mu, H\Pi_\Lambda] = 0$. Hence,

$$H\Pi_\Lambda = H\sum_\lambda \Pi^A_\lambda\Pi^B_{\Lambda\bar{\lambda}} = \sum_\lambda \Pi^A_\lambda H\Pi^B_{\Lambda\bar{\lambda}}. \tag{D.75}$$

  By taking the partial transpose of the equation above, we get

$$(H\Pi_\Lambda)^{T_A} = \sum_\lambda H^{T_A}\Pi^A_\lambda\Pi^B_{\Lambda\bar{\lambda}} = H^{T_A}\Pi_\Lambda, \tag{D.76}$$

  which, upon right multiplication with $(\mathbb{1} - \Pi_\Lambda)$, results in zero.

□

Note that the negation of the statements (i) and (iii) are the NC and EC, respectively, thereby proving their equivalence.

D.8 Entanglement negativity of the thermal cluster chain state

We analytically compute the entanglement negativity of the Gibbs and canonical ensembles of the one-dimensional cluster chain Hamiltonian. This model was chosen for both its analytical simplicity, and for the Gibbs ensemble showing non-trivial SDOE, that is, it is entangled at low temperatures, and becomes separable at high temperatures. This last property makes the cluster chain a more realistic example of the generic case, and is not shared with the classical Ising model studied in [109].

We confirm that the Gibbs state of the cluster chain exhibits sudden death of entanglement negativity, while the negativity of the canonical ensemble persists for arbitrarily high temperatures. Moreover, we find that there is no bound entanglement for the Gibbs ensemble.

D.8.1 Gibbs ensemble

Consider a connected region A with size $|A| \geq 2$ and its complement $A^C = B$ with size $|B| \geq 2$, so that the total system size N is even. Region A has two endpoints: a_L and a_R, and they are adjacent to the endpoints b_L and b_R of B, respectively. Define $A_{\text{int}} := A \setminus \{a_L, a_R\}$ to be the interior of region A, with endpoints a_L and a_R removed, and similarly with B_{int}.

To simplify the entanglement calculation, we will conjugate the cluster chain Hamiltonian $H_{\text{cc}} = -\sum_i Z_{i-1}X_iZ_{i+1}$ by CZ_{int}, defined as

$$CZ_{\text{int}} := CZ_{a_L,b_L}CZ_{a_R,b_R}\prod_{\langle i,j\rangle} CZ_{i,j} = \prod_{\substack{\langle i,j\rangle \neq \\ \langle a_L,b_L\rangle, \\ \langle a_R,b_R\rangle}} CZ_{i,j}. \tag{D.77}$$

Since CZ_{int} is a local unitary to the $A|B$ bipartition, it preserves the entanglement value of the thermal state. This results in the trivial paramagnetic Hamiltonian everywhere except on the two boundary links:

$$H'_{cc} := CZ_{\text{int}}H_{cc}CZ_{\text{int}} = H'_L + H'_R + H'_{pm}, \tag{D.78}$$

where

$$H'_L = -X_{a_L}Z_{b_L} - Z_{a_L}X_{b_L}, \tag{D.79}$$
$$H'_R = -X_{a_R}Z_{b_R} - Z_{a_R}X_{b_R}, \tag{D.80}$$
$$H'_{pm} = -\sum_{i\in A_{\text{int}}} X_i - \sum_{i\in B_{\text{int}}} X_i. \tag{D.81}$$

The Gibbs state $\rho'_\beta = e^{-\beta H'_{cc}}/\mathcal{Z}_\beta$ is, then,

$$\begin{aligned}\rho'_\beta = (e^{-\beta H'_{pm}}/\mathcal{Z}^{pm}_\beta) &\otimes (\mathbb{1} + \lambda X_{a_L}Z_{b_L} + \lambda Z_{a_L}X_{b_L} + \lambda^2 Y_{a_L}Y_{b_L}) \\ &\otimes (\mathbb{1} + \lambda X_{a_R}Z_{b_R} + \lambda Z_{a_R}X_{b_R} + \lambda^2 Y_{a_R}Y_{b_R})/2^4,\end{aligned} \tag{D.82}$$
$$\mathcal{Z}^{pm}_\beta = (2\cosh(\beta))^{N-4}, \tag{D.83}$$

where $\lambda := \tanh(\beta)$ and N is the system size. Its partial transpose is, then

$$\begin{aligned}(\rho'_\beta)^{T_A} = (e^{-\beta H'_{pm}}/\mathcal{Z}^{pm}_\beta) &\otimes (\mathbb{1} + \lambda X_{a_L}Z_{b_L} + \lambda Z_{a_L}X_{b_L} - \lambda^2 Y_{a_L}Y_{b_L}) \\ &\otimes (\mathbb{1} + \lambda X_{a_R}Z_{b_R} + \lambda Z_{a_R}X_{b_R} - \lambda^2 Y_{a_R}Y_{b_R})/2^4,\end{aligned} \tag{D.84}$$

and the logarithmic negativity $E_N = \log \|\cdot\|_1$ is

$$E_N(\rho_\beta) = E_N(\rho'_\beta) = \log\left(\left\|\mathbb{1} + \lambda X_{a_L} Z_{b_L} + \lambda Z_{a_L} X_{b_L} - \lambda^2 Y_{a_L} Y_{b_L}\right\|_1^2 / 2^4\right) \tag{D.85}$$

$$= 2\log\Big(\sum_{a,b=\pm 1} |1 + \lambda(a+b) - \lambda^2 ab|/4\Big). \tag{D.86}$$

As shown in the red curve of Fig. 5.2, the entanglement negativity is zero for small $\lambda \leq \lambda_c$. Indeed, it can be easily deduced from the expression above that $\lambda_c = \sqrt{2} - 1$. Moreover, we can further argue that the state is exactly separable for $\lambda \leq \lambda_c$ by showing an explicit separable decomposition for it.

Consider the density matrix ρ_{AB} defined by

$$\rho_\beta \propto e^{-\frac{\beta}{2} H_A}\, e^{-\frac{\beta}{2} H_B}\, \rho_{AB}\, e^{-\frac{\beta}{2} H_B}\, e^{-\frac{\beta}{2} H_A}, \tag{D.87}$$

where H_A contains the terms of H entirely supported on region A, and similarly with H_{AB}. If ρ_{AB} can be decomposed into an ensemble of states that are separable between A and B, then a separable decomposition for the full density matrix ρ_β can also be constructed.

In the case of the cluster chain Hamiltonian, and expressing ρ_{AB} only at the four closest sites near the rightmost boundary between A and B for simplicity, we have

$$\rho_{AB} = \frac{1}{2^4}\left(\mathbb{1} + \lambda Z_{a_R - 1} X_{a_R} Z_{b_R}\right)\left(\mathbb{1} + \lambda Z_{a_R} X_{b_R} Z_{b_R + 1}\right). \tag{D.88}$$

It is easy to check that $\rho_{AB} = p_1\rho_1 + p_2\rho_2 + p_3\rho_3$, with

$$\rho_1 = \frac{1}{2^4}\left(\mathbb{1} + (\lambda^2 + 2\lambda) Z_{a_R - 1} X_{a_R} Z_{b_R}\right), \tag{D.89}$$

$$\rho_2 = \frac{1}{2^4}\left(\mathbb{1} + (\lambda^2 + 2\lambda) Z_{a_R} X_{b_R} Z_{b_R + 1}\right), \tag{D.90}$$

$$\rho_3 = \frac{1}{2^4}\left(\mathbb{1} + (\lambda^2 + 2\lambda) Z_{a_R - 1} X_{a_R} Z_{b_R} \cdot Z_{a_R} X_{b_R} Z_{b_R + 1}\right), \tag{D.91}$$

and probabilities $p_1 = p_2 = \lambda/(\lambda^2 + 2\lambda)$ and $p_3 = \lambda^2/(\lambda^2 + 2\lambda)$.

Each individual ρ_i is in fact a fully separable matrix since it has the form $\mathbb{1} + \alpha\hat{P}$, where $\hat{P}$ is a Pauli operator. However, ρ_i is positive-semidefinite if and only if:

$$-1 \leq \lambda^2 + 2\lambda \leq 1. \tag{D.92}$$

Therefore the state is separable between left and right if and only if $\lambda \leq \lambda_c = \sqrt{2} - 1$.

D.8.2 Canonical ensemble

Here, we will repeat the calculation for the canonical ensemble of the 1d cluster chain. Being a nontrivial SPT, the cluster chain exhibit edge modes that become entangled when global strong symmetry is imposed. Since we are interested in the bulk entanglement, we will disregard this contribution by considering periodic boundary conditions, forming a 1d ring. Hence, a nontrivial connected region A will have two boundary points. Furthermore, contrary to the 1d Ising model, the stabilizers of the cluster chain and the symmetry operators are not independent. As we will see, this correlates the degrees of freedom inside the bulk with the ones at the boundary, and also between the boundary points.

Even though the cluster chain has $\mathbb{Z}_2 \times \mathbb{Z}_2$ symmetry generated by $\prod_i X_{2i}$ and $\prod_i X_{2i-1}$, we will only impose the strong symmetry of the diagonal component $P := \prod_i X_i = \Lambda \in \{\pm 1\}$ with projector $\Pi_\Lambda = (\mathbb{1} + \Lambda P)/2$, for simplicity. After conjugation by CZ_{int}, the symmetry becomes $P' = P_{\text{int}} P_{LR}$, where $P_{\text{int}} = \prod_{i \in A_{\text{int}}} X_i \prod_{i \in B_{\text{int}}} X_i$ and $P_{LR} = Y_{a_L} Y_{b_L} Y_{a_R} Y_{b_R}$. The strongly symmetric thermal state $\rho'_{\beta,\Lambda} = e^{-\beta H'} \Pi'_\Lambda / \mathcal{Z}_{\beta,\Lambda}$ is

$$\rho'_{\beta,\Lambda} = \Pi'_\Lambda \cdot (e^{-\beta H_{pm}} / \mathcal{Z}^{pm'}_{\beta,\Lambda}) \otimes (\mathbb{1} + \lambda X_{a_L} Z_{b_L} + \lambda Z_{a_L} X_{b_L} + \lambda^2 Y_{a_L} Y_{b_L}) \otimes (L \leftrightarrow R)/2^4, \text{ where} \tag{D.93}$$

$$\mathcal{Z}^{pm'}_{\beta,\Lambda} = (2\cosh(\beta))^{N-4}(1 + \Lambda \lambda^N)/2. \tag{D.94}$$

Now, $\rho'_{\beta,\Lambda}$ does not factorize into a tensor product over the interior of each region and their boundaries. To deal of this, we decompose $\Pi'_\Lambda = \Pi^{\text{int}}_{+1} \Pi^{LR}_\Lambda + \Pi^{\text{int}}_{-1} \Pi^{LR}_{-\Lambda}$. Then,

$$\rho'_{\beta,\Lambda} = \sum_{\mu=\pm 1} r_{N,\mu} \rho^{pm}_{\beta,\mu} \otimes \Pi^{LR}_{\mu\Lambda} (\mathbb{1} + \lambda X_{a_L} Z_{b_L} + \lambda Z_{a_L} X_{b_L} + \lambda^2 Y_{a_L} Y_{b_L}) \otimes (L \leftrightarrow R)/2^4, \text{ where} \tag{D.95}$$

$$\rho^{pm}_{\beta,\mu} = \Pi^{\text{int}}_\mu e^{-\beta H_{pm}} / \mathcal{Z}^{pm}_{\beta,\mu}, \tag{D.96}$$

$$\mathcal{Z}^{pm}_{\beta,\mu} = (2\cosh(\beta))^{N-4}(1 + \mu \lambda^{N-4})/2, \tag{D.97}$$

$$r_{N,\mu} = \mathcal{Z}^{pm}_{\beta,\mu} / \mathcal{Z}^{pm'}_{\beta,\Lambda} = (1 + \mu \lambda^{N-4})/(1 + \Lambda \lambda^N). \tag{D.98}$$

The ratio $r_{N,\mu}$ converges to 1 in the thermodynamic limit $N \to \infty$ if $\lambda < 1$ ($\beta < \infty$). Hence, we will set it to 1 in what follows.

The partial transpose is

$$(\rho'_{\beta,\Lambda})^{T_A} = \sum_{\mu=\pm 1} \rho^{pm}_{\beta,\mu} \otimes [(\mathbb{1} + \lambda X_{a_L} Z_{b_L} + \lambda Z_{a_L} X_{b_L} - \lambda^2 Y_{a_L} Y_{b_L}) \otimes (L \leftrightarrow R) + \tag{D.99}$$

$$\mu\Lambda(\mathbb{1} - \lambda X_{a_L} Z_{b_L} - \lambda Z_{a_L} X_{b_L} - \lambda^2 Y_{a_L} Y_{b_L}) \otimes (L \leftrightarrow R) P_{LR}]/2^5, \tag{D.100}$$

so the logarithmic negativity is

$$E_N(\rho_{\beta,\Lambda}) = \log\Big(\sum_{\mu=\pm1}\sum_{\substack{a_1,b_1,\\ a_2,b_2=\pm1}}\Big|\prod_{i=1,2}(1+\lambda(a_i+b_i)-\lambda^2 a_i b_i)+\mu a_1 b_1 a_2 b_2 \\ \prod_{i=1,2}(1-\lambda(a_i+b_i)-\lambda^2 a_i b_i)\Big|/2^5\Big). \tag{D.101}$$

This result is plotted on the blue curve of Fig. 5.2, which shows the persistency of entanglement for the canonical ensemble for arbitrarily small β. Also note that the final expression is independent of the global charge Λ. For finite size, however, the entanglement will depend on Λ.

The analytical results for the entanglement negativity were validated with small size ($4 \leq N \leq 6$) numerical calculations.

D.9 Persistence of fermionic negativity in the canonical ensemble

Consider the *canonical ensemble* for fermions with strong fermion parity, $\rho_{\beta,\Lambda} \propto e^{-\beta H}\Pi_\Lambda$, where $\Pi_\Lambda = \frac{1}{2}(\mathbb{1} + \Lambda P)$ projects onto a fixed fermion parity sector with $\Lambda = \pm 1$. Assume $\rho_{\beta,\Lambda}$ has zero fermionic negativity, i.e. $||\rho_{\beta,\Lambda}^{\mathcal{T}_A}||_1 = 1$, for a sequence of β converging to 0. This, combined with the fact that $\mathrm{Tr}\{\rho_{\beta,\Lambda}^{\mathcal{T}_A}\} = \mathrm{Tr}\{\rho_{\beta,\Lambda}\} = 1$, imply that $\rho_{\beta,\Lambda}$ is Hermitian and positive semi-definite [143].

Under the fermionic partial transpose, $P^{\mathcal{T}_A} = \eta_A P$, with $\eta_A = (-1)^{|A| \bmod 2}$, implying $\Pi_\Lambda \Pi_{-\eta_A\Lambda}^{\mathcal{T}_A} = \Pi_\Lambda \Pi_{-\Lambda} = 0$. Thus, the matrix $\Pi_{-\eta_A\Lambda}\rho_{\beta,\Lambda}^{\mathcal{T}_A}\Pi_{-\eta_A\Lambda}$, which is positive semidefinite, is also traceless

$$\mathrm{Tr}\big(\rho_{\beta,\Lambda}^{\mathcal{T}_A}\Pi_{-\eta_A\Lambda}\big) \propto \mathrm{Tr}\big(e^{-\beta H}\Pi_\Lambda \Pi_{-\eta_A\Lambda}^{\mathcal{T}_A}\big) = 0,$$

and thus identically zero. Since the matrix is holomorphic function of β, its derivative at $\beta = 0$ must also vanish. Negating the chain of arguments, we find that E_N^f persists if the derivative at $\beta = 0$ is nonvanishing:

$$\Pi_{-\eta_A\Lambda}(H\Pi_\Lambda)^{\mathcal{T}_A}\Pi_{-\eta_A\Lambda} \neq 0. \tag{D.102}$$

This is the fermionic analogue of the NC, which is also generically satisfied. The two-Majorana term $V = ic_i c_j$, with $c_i \in A$ and $c_j \in B$, obeys $(V\Pi_\Lambda)^{\mathcal{T}_A} = iV\Pi_{-\eta_A\Lambda}$, yielding

$$\Pi_{-\eta_A\Lambda}(V\Pi_\Lambda)^{\mathcal{T}_A}\Pi_{-\eta_A\Lambda} = iV\Pi_{-\eta_A\Lambda} \neq 0,$$

thereby providing an explicit entangling perturbation.

Appendix E

Supplementary Material for Chapter 6

E.1 Preliminaries

E.1.1 Hilbert spaces, norms, and a convexity fact:

The underlying space in our discussion is a Hilbert space, equipped with the norm induced by its inner product. For any vector $|\psi\rangle$,

$$\|\psi\| = \sqrt{\langle\psi|\psi\rangle},$$

and we call $|\psi\rangle$ normalized if $\|\psi\| = 1$.

For a linear operator A, the *operator norm* is

$$\|A\|_{\text{op}} = \sup_{\|\phi\|=1} \|A\,|\phi\rangle\,\| = \sup_{\|\phi\|=1} \sqrt{\langle\phi|\,A^\dagger A\,|\phi\rangle}. \tag{E.1}$$

When A is Hermitian, the spectral theorem provides a decomposition $A = \sum_i \lambda_i\,|v_i\rangle\langle v_i|$ with real eigenvalues λ_i. In this case

$$\|A\|_{\text{op}} = \max_i |\lambda_i|.$$

More generally, any matrix A (square or rectangular) admits a *singular value decomposition (SVD)*

$$A = U\Sigma V^\dagger,$$

where U, V are unitary (or isometries in the rectangular case) and Σ is diagonal with nonnegative entries s_i called the *singular values* of A. By definition, s_i are the eigenvalues of $|A| := \sqrt{A^\dagger A}$. This is distinct from the spectral decomposition, which applies only to normal operators. For Hermitian positive semidefinite A, the two coincide and the singular values are just the eigenvalues. In general, singular values equal the absolute values of eigenvalues only when A is normal.

The singular values provide a unified way to define matrix/operator norms, known as Schatten p-norms. For $1 \le p < \infty$,

$$\|A\|_p := \left(\operatorname{Tr}[(A^\dagger A)^{p/2}]\right)^{1/p} = \left(\sum_i |s_i|^p\right)^{1/p},$$

and for $p = \infty$,

$$\|A\|_\infty := \lim_{p\to\infty} \|A\|_p = \max_i s_i.$$

Special cases include:

- Operator norm: $\|A\|_{\text{op}} = \|A\|_\infty = \max_i s_i$,
- Trace norm: $\|A\|_1 = \sum_i |s_i|$,
- Hilbert–Schmidt norm: $\|A\|_2 = \sqrt{\sum_i s_i^2} = \sqrt{\operatorname{Tr}(A^\dagger A)}$.

These norms are always defined via singular values, which are nonnegative by construction.

A standard fact from convex analysis is that any linear functional achieves its maximum over a compact convex set at one of the extreme points of that set. In quantum theory this principle is particularly useful: for instance, when optimizing a functional of the form

$$\sigma \;\mapsto\; \langle\Psi|\,\sigma\,|\Psi\rangle\,,$$

where $|\Psi\rangle$ is fixed and σ ranges over the convex set of separable states. Since the extreme points of the separable set are precisely the pure product states $|\alpha\rangle \otimes |\beta\rangle$, the optimization reduces to this smaller class. Entangled states lie outside the separable set and thus never contribute. Consequently, the maximum overlap of a bipartite pure state with a separable state is attained on a pure product state, and is known to equal the largest Schmidt coefficient of the given state.

E.1.2 From qubits to fermions: modes, parity, and locality

Fermionic systems are most naturally described in the *second-quantized* (Fock-space) formalism, in close analogy with the bosonic framework in quantum optics. Two fermion-specific features shape locality and entanglement: (a) each mode is effectively a two-level system due to the Pauli exclusion principle, and (b) a *parity superselection rule* forbids coherent superpositions of even- and odd-parity sectors. With these ingredients, fermionic modes can be treated as qubit-like degrees of freedom in Fock space, with anticommutation relations handled automatically by the formalism—for instance, $(|00\rangle + |11\rangle)/\sqrt{2}$ plays the role of a fermionic Bell pair.

Entanglement in fermionic systems can be discussed from two complementary viewpoints: either between indistinguishable particles (first quantization) [198, 135], or between distinguishable *sets of modes* corresponding to spatial regions (second quantization) [16, 213, 72]. The latter perspective is the natural one: fermions are indistinguishable and may delocalize freely, so entanglement is defined between mode sets rather than between identical particles. Assigning modes to observers (e.g., Alice's modes in region A and Bob's in region B), correlations arise from observables on these subsystems. Unlike in bosonic or spin/qubit systems, however, observables supported on disjoint fermionic regions may anticommute, introducing algebraic constraints that make the analysis of locality and entanglement more subtle.

The translation from qubits to fermions can be consolidated into the following structural principles:

(i) Mode dimensionality. Each fermionic mode is two-dimensional, with basis $\{|0\rangle, |1\rangle\}$ for empty/occupied. At this level, a single mode is isomorphic to a qubit.

(ii) Parity superselection. Physical states and operations preserve fermion-number parity. Coherent superpositions of different parity sectors are forbidden. For one mode, $\alpha|0\rangle + \beta|1\rangle$ is not physical, whereas $p|0\rangle\langle 0| + (1-p)|1\rangle\langle 1|$ is. For two modes, $(|00\rangle + |11\rangle)/\sqrt{2}$ is allowed (even parity), while $(|00\rangle + |01\rangle)/\sqrt{2}$ is not.

(iii) Observable algebras and locality. Creation/annihilation operators satisfy the canonical anticommutation relations

$$\{\hat{c}_i, \hat{c}_j^\dagger\} = \delta_{ij}, \qquad \{\hat{c}_i, \hat{c}_j\} = \{\hat{c}_i^\dagger, \hat{c}_j^\dagger\} = 0, \tag{E.2}$$

and bare mode operators on disjoint regions do not commute. Physical measurements, however, are generated by *parity-preserving* (even) operators, and these *do* commute across spacelike-separated regions:

$$[O_A, O_B] = 0 \quad \text{for } O_A, O_B \text{ even, supported on } A, B. \tag{E.3}$$

Entanglement and separability are therefore defined with respect to the even (parity-preserving) local algebras.

(iv) Tensor products and operator ordering. Fermionic operators act on the global antisymmetric Fock space. Writing $\hat{c}_1 \otimes I$ or $I \otimes \hat{c}_2$ naively would enforce commutation rather than anticommutation. When a tensor-product representation is required (e.g., to map to spins), Jordan–Wigner or related encodings are employed to keep track of operator order and parity strings.

(v) Subsystem parity. Let $\hat{N}_A = \sum_{k=1}^{n_A} \hat{c}_k^\dagger \hat{c}_k$ be the number operator on a subsystem A with n_A modes. The fermion-number parity on A is

$$P^A = (-1)^{\hat{N}_A} = e^{i\pi \hat{N}_A} = \prod_{k=1}^{n_A} \big(1 - 2\,\hat{c}_k^\dagger \hat{c}_k\big), \tag{E.4}$$

which satisfies $(P^A)^2 = I^A$. Physical states are block-diagonal in the eigenspaces of P^A. (Equivalently, using Majorana operators introduced in Sec. E.1.3, one can write $P^A = i^{\,n_A} \prod_{k=1}^{n_A} \gamma_{2k-1}^A \gamma_{2k}^A$.)

These structural ingredients—mode dimensionality, parity superselection, algebraic locality, operator ordering, and subsystem parity—underlie the passage from qubits to fermions and govern the formulation of fermionic quantum information.

E.1.3 Majorana operators and covariance matrices

For n fermionic modes with creation/annihilation operators $a_j^\dagger, a_j$, introduce $2n$ Hermitian Majorana operators

$$\gamma_{2j-1} = \hat{c}_j + \hat{c}_j^\dagger, \qquad \gamma_{2j} = i(\hat{c}_j - \hat{c}_j^\dagger), \tag{E.5}$$

which generate a real Clifford algebra:

$$\{\gamma_p, \gamma_q\} = 2\delta_{pq}, \qquad \gamma_p^\dagger = \gamma_p. \tag{E.6}$$

Given a (mixed) state ρ, its (Majorana) CM $M \in \mathbb{R}^{2n\times 2n}$ is

$$M_{pq} = \frac{i}{2}\,\mathrm{Tr}\big(\rho[\gamma_q,\gamma_p]\big) = \begin{cases} \mathrm{Tr}(i\gamma_q\gamma_p\,\rho)\,, & q\neq p,\\ 0, & q=p. \end{cases} \tag{E.7}$$

Thus M is real and antisymmetric. Writing iM (which is Hermitian), the *bona fide* (physicality) condition for CMs is the spectral bound

$$\mathrm{spec}(iM) \subseteq [-1,1], \tag{E.8}$$

equivalently, there exists $O \in SO(2n)$ and real numbers $\lambda_1,\ldots,\lambda_n \in [-1,1]$ such that

$$M = O\bigoplus_{j=1}^{n}\begin{bmatrix} 0 & \lambda_j \\ -\lambda_j & 0\end{bmatrix}O^T. \tag{E.9}$$

In particular, the singular values of M are at most 1. For later use, note $(\gamma_p\gamma_q)^\dagger = -\gamma_p\gamma_q$ for $p\neq q$, so $(\gamma_p\gamma_q)^2 = I$ and its eigenvalues are ± 1.

E.1.4 Fermionic Gaussian states

A (parity-preserving) state ρ on n modes is called *fermionic Gaussian* if it can be expressed as the (properly normalized) exponential of a quadratic form in Majoranas:

$$\rho \;=\; C\,\exp\!\big(\tfrac{i}{2}\,\gamma^T h\,\gamma\big)\,, \tag{E.10}$$

with $h \in \mathbb{R}^{2n\times 2n}$ antisymmetric and C determined by the normalization $\mathrm{Tr}\,\rho = 1$. Equivalently, ρ is Gaussian if and only if all higher-order correlators are fixed by the two-point ones via Wick's theorem: for any even Majorana monomial $\gamma(x)$,

$$\mathrm{Tr}\big(i^{|x|/2}\gamma(x)\,\rho\big) \;=\; \mathrm{Pf}\big(M[x]\big), \tag{E.11}$$

where $M[x]$ is the principal submatrix of the CM M supported on the indices of x. Expectations of odd-weight monomials vanish.

Purity and canonical form. For Gaussian states, the eigenvalues $\{\lambda_j\}$ in the canonical form (E.9) completely determine the spectrum of ρ. Purity is equivalent to

$$M^2 = -I \qquad\Leftrightarrow\qquad |\lambda_j| = 1 \;\;\forall j, \tag{E.12}$$

while mixed Gaussians have $|\lambda_j| < 1$ for some j. In the basis that block-diagonalizes M, the state factorizes into a product over n two-Majorana subsystems:

$$\rho \;=\; \frac{1}{2^n}\prod_{j=1}^{n}\big(I + i\lambda_j\,\tilde{\gamma}_{2j-1}\tilde{\gamma}_{2j}\big), \tag{E.13}$$

where $\tilde{\gamma} = O^T\gamma$. Using the identity

$$\exp\big(\alpha\, i\tilde{\gamma}_{2j-1}\tilde{\gamma}_{2j}\big) = \cosh\alpha \;+\; \sinh\alpha\, i\tilde{\gamma}_{2j-1}\tilde{\gamma}_{2j},$$

and comparing with (E.10) gives the matrix functional relation

$$\tanh\big(\tfrac{h}{2}\big) = M \tag{E.14}$$

(in the simultaneous block-diagonal basis of h and M).

Single-mode states. Parity superselection fixes the structure of one-mode states: the only pure states are $|0\rangle$ and $|1\rangle$, and mixed states are diagonal in this basis. Every one-mode state is Gaussian and admits the representation

$$\rho \;=\; \tfrac{1}{2}\big(I + i\lambda\,\gamma_1\gamma_2\big), \qquad \lambda \in [-1, 1], \tag{E.15}$$

with $\lambda = \pm 1$ corresponding to $|0\rangle\langle 0|$ and $|1\rangle\langle 1|$, and $\lambda = 0$ the maximally mixed state.

Vacuum and Bell-pair examples. Let $|\mathrm{vac}\rangle$ denote the fermionic Fock vacuum. For a single mode with Majoranas (γ_1, γ_2) its CM is

$$M_{|\mathrm{vac}\rangle} = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}. \tag{E.16}$$

For two modes with Majoranas $(\gamma_1, \gamma_2, \gamma_3, \gamma_4)$, the even Bell states

$$|\Phi^{\pm}\rangle = \tfrac{1}{\sqrt{2}}(|00\rangle \pm |11\rangle), \quad |\Psi^{\pm}\rangle = \tfrac{1}{\sqrt{2}}(|01\rangle \pm |10\rangle) \tag{E.17}$$

have CMs (with respect to the natural ordering) given by

$$M_{\Phi^{\pm}} = \begin{bmatrix} 0 & 0 & 0 & \pm 1 \\ 0 & 0 & \pm 1 & 0 \\ 0 & \mp 1 & 0 & 0 \\ \mp 1 & 0 & 0 & 0 \end{bmatrix}, \qquad M_{\Psi^{\pm}} = \begin{bmatrix} 0 & 0 & 0 & \mp 1 \\ 0 & 0 & \pm 1 & 0 \\ 0 & \mp 1 & 0 & 0 \\ \pm 1 & 0 & 0 & 0 \end{bmatrix}. \tag{E.18}$$

More generally, the maximally entangled m-mode state between two m-mode registers,

$$|\mathrm{EPR}\rangle \;=\; \frac{1}{\sqrt{2^m}} \sum_{j=0}^{2^m-1} |j,j\rangle\,, \tag{E.19}$$

has a CM with block off-diagonal form

$$M_{\mathrm{EPR}} = \begin{bmatrix} 0 & I_m \\ -I_m & 0 \end{bmatrix}. \tag{E.20}$$

Locality, ordering, and signs. Because $\{\hat{c}_i^\dagger, \hat{c}_j^\dagger\} = 0$, the order of creation operators is crucial:

$$\hat{c}_1^\dagger \hat{c}_2^\dagger\,|\mathrm{vac}\rangle = -\,\hat{c}_2^\dagger \hat{c}_1^\dagger\,|\mathrm{vac}\rangle\,. \tag{E.21}$$

A consistent global convention for operator ordering must be fixed, after which relative minus signs in multi-mode states follow automatically. This is a bookkeeping feature of second quantization and does not obstruct operational locality, provided we restrict to even observables.

Two structural corollaries. (i) Since M is real and antisymmetric, its eigenvalues come in pairs $\pm i\lambda_j$ with $\lambda_j \in \mathbb{R}$. The physicality condition (E.8) is equivalently $|\lambda_j| \leq 1$. (ii) Traces of odd Majorana monomials vanish.

E.1.5 Fermionic Gaussian channels

A *fermionic Gaussian channel* (FGC) is a completely positive trace-preserving (CPTP) map that sends fermionic Gaussian states to fermionic Gaussian states [78, 27]. As in the bosonic case, FGCs are fully determined by their action on CMs.

An $m \to n$ mode FGC $\mathcal{N}^{A\to B}$ is specified by a real $2n \times 2m$ matrix X and a real antisymmetric $2n \times 2n$ matrix N_B, acting as

$$M_A \;\longmapsto\; X M_A X^T + N_B, \tag{E.22}$$

for any input CM M_A. Bona fid constraints on (X, N_B) can be most easily obtained via the Choi–Jamiołkowski (Choi) state, defined by

$$C_{\mathcal{N}}^{AB} = (\mathcal{N} \otimes I_A)(\Phi_{AA'}^{+}), \tag{E.23}$$

with $\Phi^{+}_{AA'}$ a maximally entangled fermionic Gaussian Bell state. Since Gaussianity is preserved, $C^{AB}_{\mathcal{N}}$ is Gaussian with CM

$$M_{AB} = \begin{pmatrix} N_B & X \\ -X^T & 0 \end{pmatrix}. \tag{E.24}$$

The map (E.22) is a valid FGC iff M_{AB} is a bona fide CM, i.e.

$$I + iN_B - XX^T \;\geq\; 0. \tag{E.25}$$

Implying: (i) $XX^T \leq I$, and (ii) $I + iN_B \geq 0$.

Degradable, anti-degradable, and related channels

An FGC $\mathcal{N}$ is *degradable* if its complementary channel $\mathcal{N}^c$ can be simulated from $\mathcal{N}$ by a CPTP degrading map $\mathcal{D}$:

$$\mathcal{N}^c = \mathcal{D} \circ \mathcal{N}. \tag{E.26}$$

It is *anti-degradable* if there exists a CPTP map $\mathcal{A}$ such that

$$\mathcal{N} = \mathcal{A} \circ \mathcal{N}^c, \tag{E.27}$$

in which case its quantum capacity vanishes by no-cloning theorem.

A special case is the *replacement channel*, which outputs a fixed state σ regardless of input:

$$\mathcal{N}(\rho) = \sigma, \quad \forall \rho. \tag{E.28}$$

If σ is separable from any reference, $\mathcal{N}$ is *entanglement-breaking* [93], destroying all input–reference entanglement.

E.1.6 Useful lemmas:

Lemma E.1 ([20, Proposition 3])**.** *Let ρ and σ be two fermionic states with correlation matrices M_ρ and M_σ, respectively. Then, their trace distance is lower bounded by the operator norm of the difference of their correlation matrices:*

$$\|M_\rho - M_\sigma\|_\infty \leq \|\rho - \sigma\|_1. \tag{E.29}$$

Lemma E.2 ([20, Theorem 6])**.** *Let ρ and σ be two fermionic states (possibly mixed) with CMs M_ρ and M_σ, respectively. Then, the trace distance between ρ and σ is upper bounded by one half of the trace norm of the difference of their CMs:*

$$\|\rho - \sigma\|_1 \leq \frac{1}{2}\|M_\rho - M_\sigma\|_1. \tag{E.30}$$

Lemma E.3. *Given a CM M the following inequality holds:*

$$\sum_k (M_{ki})^2 \leq 1,$$

Proof. Since $M^T M \leq I$, it follows that for all i,

$$(M^T M)_{ii} = \sum_k M_{ki} M_{ki} \leq \delta_{ii}.$$

which proves the claim. □

Lemma E.4. *Let M be the CM of a fermionic Gaussian state. Then, for all i, j,*

$$|M_{ij}| \leq 1.$$

Proof. By definition, the off-diagonal entries of the CM M are given by

$$M_{ij} = \mathrm{Tr}(i\gamma_j\gamma_i\rho),$$

where γ_i and γ_j are Majorana operators satisfying $\gamma_k = \gamma_k^\dagger$, $\gamma_k^2 = I$, and $\{\gamma_i, \gamma_j\} = 2\delta_{ij}$.

Define the operator $A := i\gamma_j\gamma_i$. Then A is Hermitian, and since

$$A^2 = (i\gamma_j\gamma_i)^2 = -\gamma_j\gamma_i\gamma_j\gamma_i = -\gamma_j(\gamma_i\gamma_j)\gamma_i = -\gamma_j(-\gamma_j\gamma_i)\gamma_i = \gamma_j\gamma_j\gamma_i\gamma_i = I,$$

we have $A^2 = I$, so the eigenvalues of A are ± 1. Therefore, its operator norm is $\|A\|_{\mathrm{op}} = 1$.

Now, since $\rho = \sum_k p_k |\psi_k\rangle\langle\psi_k|$ is a valid density matrix (positive semidefinite, $\mathrm{Tr}(\rho) =$

1), the expectation value satisfies

$$\begin{aligned}
|M_{ij}| &= |\operatorname{Tr}(A\rho)| \\
&= \left| \sum_k p_k \langle\psi_k| A |\psi_k\rangle \right| \\
&\leq \sum_k p_k |\langle\psi_k| A |\psi_k\rangle| \\
&\leq \sum_k p_k . \|A\|_{\text{op}} \\
&= \|A\|_{\text{op}} \\
&= 1,
\end{aligned}$$

where the first inequality follows from Jensen's inequality applied to the absolute value function which is a convex function (it also follows more easily from triangle inequality though), and the second inequality follows from the definition of the operator norm Eq. (E.1). This proves the claim. □

Lemma E.5. *For any bipartite pure state $|\psi\rangle^{AB}$ with maximal Schmidt coefficient $\lambda_1(\psi)$, the maximal overlap with a separable state σ^{AB} is bounded by*

$$\langle\psi| \sigma |\psi\rangle \leq \lambda_1(\psi).$$

Proof. The functional $\sigma \mapsto \langle\psi| \sigma |\psi\rangle$ is linear in σ. By convex analysis, its maximum over the compact convex set of separable states

$$\sigma^{AB} = \sum_x p(x) \sigma_x^A \otimes \sigma_x^B$$

is attained at an extreme point, i.e., at a pure product state $|\alpha\rangle \otimes |\beta\rangle$. The maximum squared overlap of $|\psi\rangle$ with a product state is known to equal the largest Schmidt coefficient $\lambda_1(\psi)$. Therefore,

$$\max_{\text{sep } \sigma} \langle\psi| \sigma |\psi\rangle \leq \lambda_1(\psi).$$

□

E.2 Extendibility of Gaussian fermions

Let ρ^{AB} be a fermionic Gaussian state with CM

$$M_{AB} = \begin{pmatrix} M_A & X \\ -X^T & M_B \end{pmatrix}.$$

By definition, M_{AB} is real and antisymmetric. The diagonal blocks M_A and M_B are the reduced CMs of the subsystems A and B, hence are real and antisymmetric themselves. The off-diagonal block is otherwise unconstrained apart from the requirement that the full CM remain real and antisymmetric; accordingly, we may take X to be an arbitrary real matrix (in particular, X need not be antisymmetric).

The set of Gaussian states is closed under taking extensions by Theorem 1 of the main text; in particular, a (k_1, k_2)-extension of ρ^{AB} exists within the Gaussian family and has a CM of the form given in Eq. (E.31). Whenever that extended CM is valid, all moments (including higher-order ones) are finite; see, e.g., Lemma E.4. The crucial observation, however, is that the additional blocks Y and Z must themselves be antisymmetric, a property that is far from obvious a priori. Recognizing and proving this fact is essential in the fermionic setting. In the bosonic case, the corresponding statement was established in [122] by combining extendibility with structural properties of the bosonic CM. Our approach takes a different route: the result is obtained directly from the fermionic algebra of Majorana operators together with the definition of extendibility. This highlights that the fermionic case cannot simply be reduced to the bosonic one, but requires its own algebraic foundations.

Lemma E.6. *The matrices Y and Z in Eq. (E.31) is antisymmetric.*

Proof. Consider two distinct copies B_m and B_n within the B-part of a symmetric (k_1, k_2)-extension of ρ^{AB}. Let $\gamma_q^{B_m}$ and $\gamma_p^{B_n}$ denote Majorana operators acting on B_m and B_n, respectively. Writing (by slight abuse of notation) $M_{(qp)} \equiv Y_{qp}$ for the (q,p)-entry of the CM block coupling B_m to B_n, we have

$$\begin{aligned}
Y_{qp} &= i\ \mathrm{Tr}\big(\gamma_q^{B_m}\gamma_p^{B_n}\ \rho^{A_1\ldots A_{k_1},B_1\ldots B_{k_2}}\big) \\
&= \mathrm{Tr}\big(i\,\gamma_q^{B_m}\gamma_p^{B_n}\,U_{mn}\,\widetilde{\rho}^{A_1\ldots A_{k_1},B_1\ldots B_{k_2}}\,U_{mn}^\dagger\big) \\
&= \mathrm{Tr}\big(i\,U_{mn}^\dagger\gamma_q^{B_m}U_{mn}\,U_{mn}^\dagger\gamma_p^{B_n}U_{mn}\,\widetilde{\rho}^{A_1\ldots A_{k_1},B_1\ldots B_{k_2}}\big) \\
&= \mathrm{Tr}\big(i\,\gamma_q^{B_n}\gamma_p^{B_m}\,\widetilde{\rho}^{A_1\ldots A_{k_1},B_1\ldots B_{k_2}}\big) \\
&= -\ \mathrm{Tr}\big(i\,\gamma_p^{B_m}\gamma_q^{B_n}\,\widetilde{\rho}^{A_1\ldots A_{k_1},B_1\ldots B_{k_2}}\big) \\
&= -\,Y_{pq}.
\end{aligned}$$

Here U_{mn} is the unitary that swaps the subsystems B_m and B_n (an element of the permutation representation of S_{k_2}). The second line uses symmetric extendibility; the third uses cyclicity of the trace; the fourth uses $U_{mn}^\dagger \gamma_q^{B_m} U_{mn} = \gamma_q^{B_n}$ and $U_{mn}^\dagger \gamma_p^{B_n} U_{mn} = \gamma_p^{B_m}$; and the penultimate equality follows from the anticommutation of Majorana operators acting on disjoint subsystems. Hence $Y_{qp} = -Y_{pq}$, i.e., Y is antisymmetric. The same reasoning applies to Z, which is therefore antisymmetric as well.

□

The permutation-invariance requirement in the definition of symmetric extendibility can in fact be eliminated. One may define a (k_1, k_2)-extension of ρ^{AB} simply as a state $\widetilde{\rho}^{A_1 \ldots A_{k_1}, B_1 \ldots B_{k_2}}$ whose two-body marginals satisfy $\rho^{A_i B_j} = \rho^{AB}$ for all pairs (i, j), without any symmetry assumption. With this weaker definition, extendibility and symmetric extendibility turn out to be equivalent. The implication from symmetric to general is trivial, while the converse follows by averaging: given any extension $\widetilde{\rho}^{A_1 \ldots A_{k_1}, B_1 \ldots B_{k_2}}$, one obtains a symmetric extension by twirling it over the product of permutation groups $S_{k_1} \times S_{k_2}$,

$$\widetilde{\rho}^{A_1 \ldots A_{k_1}, B_1 \ldots B_{k_2}} = \frac{1}{k_1!\, k_2!} \sum_{\pi_a \in S_{k_1}} \sum_{\pi_b \in S_{k_2}} \left(U_{\pi_a}^{A_1 \ldots A_{k_1}} \otimes U_{\pi_b}^{B_1 \ldots B_{k_2}} \right) \widetilde{\rho}^{A_1 \ldots A_{k_1}, B_1 \ldots B_{k_2}} \left(U_{\pi_a}^{A_1 \ldots A_{k_1}} \otimes U_{\pi_b}^{B_1 \ldots B_{k_2}} \right)^\dagger,$$

where $U_{\pi_a}^{A_1 \ldots A_{k_1}}$ and $U_{\pi_b}^{B_1 \ldots B_{k_2}}$ implement the permutations π_a and π_b.

As a consequence, for a Gaussian bipartite state ρ_{G}^{AB} it makes no difference whether one demands Gaussianity, symmetry, or both: the existence of a Gaussian symmetric (k_1, k_2)-extension is equivalent to the existence of a Gaussian extension without symmetry, and both are in turn equivalent to the existence of a symmetric or general (k_1, k_2)-extension that need not be Gaussian.

E.3 Proof of Theorem 6.2

Let $\rho^{A_1\ldots A_{k_1}B_1\ldots B_{k_2}}$ be a symmetric (k_1,k_2)-extension of the bipartite Gaussian state ρ_{G}^{AB}. Even if $\widetilde{\rho}$ is not Gaussian, its second moments define a CM of the form

$$M_{A_1\ldots A_{k_1}B_1\ldots B_{k_2}} = \begin{pmatrix} M_{AA} & \widetilde{M}_{AB} \\ -M_{AB}^T & M_{BB} \end{pmatrix}_{2\times 2}$$
$$= \begin{pmatrix} M_A & Z & \cdots & Z & X & X & \cdots & X \\ Z & M_A & \cdots & \vdots & X & X & \cdots & X \\ \vdots & \vdots & \ddots & Z & \vdots & \vdots & \ddots & \vdots \\ Z & \cdots & Z & M_A & X & X & \cdots & X \\ -X^T & -X^T & \cdots & -X^T & M_B & Y & \cdots & Y \\ -X^T & -X^T & \cdots & -X^T & Y & M_B & \cdots & Y \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots \\ -X^T & -X^T & \cdots & -X^T & Y & Y & \cdots & M_B \end{pmatrix}_{(k_1+k_2)\times(k_1+k_2)}, \tag{E.31}$$

where permutation symmetry of the extension enforces the block-constant structure

$$M_{AA} = \begin{pmatrix} M_A & Z & \ldots & Z \\ Z & M_A & \ldots & \vdots \\ \vdots & \vdots & \ddots & Z \\ Z & \ldots & Z & M_A \end{pmatrix}_{k_1\times k_1}, \quad M_{BB} = \begin{pmatrix} M_B & Y & \ldots & Y \\ Y & M_B & \ldots & \vdots \\ \vdots & \vdots & \ddots & Y \\ Y & \ldots & Y & M_B \end{pmatrix}_{k_2\times k_2}, \tag{E.32}$$

$$M_{AB} = \begin{pmatrix} X & X & \ldots & X \\ X & X & \ldots & \vdots \\ \vdots & \vdots & \ddots & X \\ X & \ldots & X & X \end{pmatrix}_{k_1\times k_2}. \tag{E.33}$$

Here $X \in \mathbb{R}^{2n_A\times 2n_B}$ is arbitrary, while $Y \in \mathbb{R}^{2n_B\times 2n_B}$ and $Z \in \mathbb{R}^{2n_A\times 2n_A}$ are required to be real antisymmetric matrices. physically, X, Y, Z encode inter-mode correlations across distinct fermionic subsystems.

An extension with CM M is physical iff the fermionic physicality condition holds:

$$I + i\widetilde{M} = \begin{pmatrix} I + iM_{AA} & iM_{AB} \\ -iM_{AB}^T & I + iM_{BB} \end{pmatrix} \geq 0. \tag{E.34}$$

To analyse (E.34) we use the standard Schur-complement characterization of positivity of Hermitian block matrices [92, Theorem 1.12]:

$$\begin{pmatrix} A & C \\ C^\dagger & B \end{pmatrix} \geq 0 \quad \Leftrightarrow \quad A \geq 0, \quad B \geq C^\dagger A^{-1} C,$$

with the non-invertible case handled by the usual limiting argument [122]. The same equivalence holds with A and B interchanged. It is useful to write Eq. (E.34) as follows:

$$I + iM = \begin{pmatrix} I + iM_{AA} & iM_{AB} \\ (iM_{AB})^\dagger & I + iM_{BB} \end{pmatrix} \geq 0.$$

The Schur complement then reduces to the following condition:

$$I + iM_{AA} \geq 0, \tag{E.35}$$

$$I + iM_{BB} - M_{AB}^\dagger (I + iM_{AA})^{-1} M_{AB} \geq 0. \tag{E.36}$$

(These in particular imply $I + iM_{BB} \geq 0$ as remarked earlier.)

It is convenient to decompose the block-constant matrices into the symmetric direction (span of the uniform vector) and its orthogonal complement. Define

$$|+\rangle_k \coloneqq \frac{1}{\sqrt{k}} \sum_{j=1}^{k} |j\rangle$$

A direct expansion yields the useful identities (we keep the explicit intermediate step because it is instructive):

$$\begin{aligned} I + iM_{AA} &= \begin{pmatrix} I + iM_A & iZ & \cdots & iZ \\ iZ & I + iM_A & \cdots & \vdots \\ \vdots & \vdots & \ddots & iZ \\ iZ & \cdots & iZ & I + iM_A \end{pmatrix}_{2k_1 \times 2k_1} \\ &= k_1 \, |+\rangle\langle+|_{k_1} \otimes iZ + I_{k_1} \otimes (I + iM_A - iZ) \qquad \text{(E.37)} \\ &= |+\rangle\langle+|_{k_1} \otimes \big(I + iM_A + (k_1 - 1)iZ\big) + \big(I_{k_1} - |+\rangle\langle+|_{k_1}\big) \otimes \big(I + iM_A - iZ\big). \qquad \text{(E.38)} \end{aligned}$$

The analogous decomposition on the B side is

$$I + iM_{BB} = |+\rangle\langle+|_{k_2} \otimes \big(I + iM_B + (k_2 - 1)iY\big) + \big(I_{k_2} - |+\rangle\langle+|_{k_2}\big) \otimes \big(I + iM_B - iY\big). \tag{E.39}$$

Finally the off-diagonal block is proportional to the cross projector:

$$iM_{AB} = \sqrt{k_1 k_2}\ |+\rangle_{k_1}\langle +|_{k_2} \otimes iX. \tag{E.40}$$

Because $|+\rangle\langle+|$ and $I - |+\rangle\langle+|$ are orthogonal projectors, the large matrices above decompose into direct sums acting on orthogonal subspaces. Consequently the Schur conditions (E.35)–(E.36) reduce to the following four matrix inequalities, obtained by imposing positivity separately on the symmetric sector and on its orthogonal complement:

$$I + iM_A + (k_1 - 1)iZ \geq 0, \tag{E.41}$$

$$I + iM_A - iZ \geq 0, \tag{E.42}$$

$$I + iM_B + (k_2 - 1)iY - k_1 k_2\, X^T \big(I + iM_A + (k_1 - 1)iZ\big)^{-1} X \geq 0, \tag{E.43}$$

$$I + iM_B - iY \geq 0. \tag{E.44}$$

Thus (E.41)–(E.44) are necessary conditions for the CM of a (k_1, k_2)-extension to be physical. Moreover, when the global state is Gaussian these conditions are also sufficient (second moments completely characterize Gaussian states).

These condition can be succinctly imposed by asking the following matrix to satisfy the fermionic bona fide condition:

$$\begin{pmatrix} (M_A - Z) & 0 & 0 & 0 \\ 0 & M_A + (k_1 - 1)Z & \sqrt{k_1 k_2}X & 0 \\ 0 & -\sqrt{k_1 k_2}X^T & M_B + (k_2 - 1)Y & 0 \\ 0 & 0 & 0 & (M_B - Y) \end{pmatrix} \tag{E.45}$$

Define $\Delta_A = M_A - Z$ and $\Delta_B = M_B - Y$, and we rewrite the matrix as follows:

$$\begin{pmatrix} \Delta_A & 0 & 0 & 0 \\ 0 & k_1 M_A - (k_1 - 1)\Delta_A & \sqrt{k_1 k_2}X & 0 \\ 0 & -\sqrt{k_1 k_2}X^T & k_2 M_B - (k_2 - 1)\Delta_B & 0 \\ 0 & 0 & 0 & \Delta_B \end{pmatrix} \tag{E.46}$$

Hence, a necessary condition for (k_1, k_2)-extendibility of any fermionic state is existence of valid CMs Δ_A and Δ_B such that the following is a valid CM:

$$\begin{pmatrix} k_1 M_A - (k_1 - 1)\Delta_A & \sqrt{k_1 k_2} X \\ -\sqrt{k_1 k_2} X^T & k_2 M_B - (k_2 - 1)\Delta_B \end{pmatrix} \tag{E.47}$$

When ρ^{AB} is Gaussian, second moments are sufficient to characterize extendibility, and the condition above becomes necessary and sufficient. □

E.4 Proof of Lemma 6.3 of main text

Proof.— Assume the Gaussian state is (k_1, k_2) extendible. Then four inequalities must be satisfied. Set

$$A := I + iM_A + (k_1 - 1)\, iZ, \qquad B := I + iM_B + (k_2 - 1)\, iY.$$

The constraints (E.41), (E.43) read

$$A \geq 0, \tag{E.41}$$
$$B - k_1 k_2\, X^T A^{-1} X \geq 0, \tag{E.43}$$

Since M_A, Z, M_B, Y are real antisymmetric, the matrices iM_A, iZ, iM_B, iY are Hermitian with $(iM_A)^T = -iM_A$, $(iZ)^T = -iZ$, etc. Hence

$$A^T = 2I - A, \qquad B^T = 2I - B.$$

Spectrum is invariant under transpose, so $A \geq 0 \Rightarrow A^T \geq 0$, and thus

$$0 \;\leq\; A \;\leq\; 2I. \tag{E.48}$$

As A is positive Hermitian, (E.48) implies $\lambda_{\max}(A) \leq 2$, and therefore $\lambda_{\min}(A^{-1}) = (\lambda_{\max}(A))^{-1} \geq \frac{1}{2}$, i.e.,

$$A^{-1} \;\geq\; \tfrac{1}{2} I. \tag{E.49}$$

From (E.43) and $A^{-1} \geq 0$ we obtain $B \geq 0$. The same transpose argument then gives

$$0 \;\leq\; B \;\leq\; 2I. \tag{E.50}$$

Combining (E.50) with (E.43),

$$2I \;\geq\; B \;\geq\; k_1 k_2\, X^T A^{-1} X. \tag{E.51}$$

By (E.49). Since $C \geq D$ implies $E^\dagger C E \geq E^\dagger D E$ for any E, taking $C = A^{-1}$, $D = \frac{1}{2} I$, and $E = X$ yields

$$X^T A^{-1} X \;\geq\; \tfrac{1}{2}\, X^T X. \tag{E.52}$$

Substituting (E.52) into (E.51) gives

$$2I \;\geq\; k_1 k_2\, X^T A^{-1} X \;\geq\; \tfrac{k_1 k_2}{2}\, X^T X. \tag{E.53}$$

Finally, multiplying by the positive scalar $\frac{2}{k_1 k_2}$ preserves order and gives

$$\frac{4\, I}{k_1 k_2} \;\geq\; X^T X. \tag{E.54}$$

E.5 Separability in Fermionic Gaussian States and Theorem 6.5

In fermionic systems, the notion of a product state is more subtle than in qubit systems due to anti-commutation relations between operators on different subsystems [16]. We briefly distill the relevant notions from [16]. As an example bipartite fermionic state ρ^{AB} is a *product state* if it can be written as $\rho^{AB} = \rho^A \otimes \rho^B$. Following [16], we denote the set of such states by $\mathcal{P}2$. Here, neither the marginals ρ^A, ρ^B nor the global state ρ^{AB} are required to be *physical*, i.e., they may fail to commute with the parity operator. If we further require these states to be physical, the resulting set is denoted by $\mathcal{P}2_\pi$. A different definition arises from factorization of expectation values: a fermionic state ρ^{AB} is a product if, for all (possibly non-physical and non-commuting) observables O^A and O^B,

$$\mathrm{Tr}\big[\rho^{AB}(O^A O^B)\big] \;=\; \mathrm{Tr}\big(\rho^A O^A\big)\; \mathrm{Tr}\big(\rho^B O^B\big). \tag{E.55}$$

Let $\mathcal{P}3$ denote the set of states satisfying this factorization. In general, $\mathcal{P}2 \neq \mathcal{P}3$. If, however, we restrict the *states* (not the observables) to be physical, we obtain $\mathcal{P}3_\pi$, and the crucial identity holds:

$$\mathcal{P}2_\pi \;=\; \mathcal{P}3_\pi. \tag{E.56}$$

The proof is given in [16, App. A.6] (see also [157]). For intuition, we sketch an analogous argument in the qubit setting.

Qubit analogy. In qubit systems, a bipartite state ρ^{AB} is product iff

$$\mathrm{Tr}\big[\rho^{AB}(O_A \otimes O_B)\big] \;=\; \mathrm{Tr}\big(\rho^A O_A\big)\;\mathrm{Tr}\big(\rho^B O_B\big) \quad \forall\, O_A, O_B. \tag{E.57}$$

One direction is immediate. For the converse, recall that $\rho = \sigma$ iff $\mathrm{Tr}(\rho O) = \mathrm{Tr}(\sigma O)$ for all O. Since any O^{AB} is a linear combination of product operators $O_A \otimes O_B$, Eq. (E.57) is equivalent to

$$\mathrm{Tr}\big[\rho^{AB} O^{AB}\big] \;=\; \mathrm{Tr}\big(\rho^A O^A\big)\;\mathrm{Tr}\big(\rho^B O^B\big) \quad \forall\, O^{AB}, \tag{E.58}$$

and (E.58) implies (E.57) trivially. This is the essence of the equivalence $\mathcal{P}2_\pi = \mathcal{P}3_\pi$.

Separable sets. There are two natural separability constructions. One may take convex combinations of $\mathcal{P}2$ states and require the mixture to be physical; denote this by $\mathcal{S}2'_\pi$. Alternatively, one may take convex combinations of $\mathcal{P}2_\pi (= \mathcal{P}3_\pi)$; denote this by $\mathcal{S}2_\pi$. In $\mathcal{S}2'_\pi$, the components may be unphysical but the mixture is physical, whereas in $\mathcal{S}2_\pi$ each component is already physical. Clearly, $\mathcal{S}2'_\pi \supset \mathcal{S}2_\pi$.

For *Gaussian* fermionic states, separability relative to the smaller set $\mathcal{S}2_\pi$ exhibits a striking feature: Gaussian separable states are *quantum-classical binding*, in the sense that classical correlations cannot appear without quantum correlations. In particular, one cannot generate classical correlations by Gaussian LOCC. One proof uses Theorem 5; another follows directly from

$$G\mathcal{S}3_\pi \;=\; G\mathcal{P}2_\pi \;=\; G\mathcal{P}3_\pi, \tag{E.59}$$

which precludes classical correlations in the absence of entanglement under Gaussian operations and classical communication.

We now present the proof of Theorem 5, establishing that, for fermionic Gaussian states, classical correlations are present if and only if entanglement is present. This theorem, together with the extendibility framework, indicates that $\mathcal{S}2_\pi$ is the natural candidate for defining separable fermionic Gaussian states. Furthermore, it lends support to a more general picture: even beyond the Gaussian setting, extendibility requires the second-order moments to vanish, consistent with the intuition behind this definition.

Proof of Theorem 5. We combine the fermionic superselection rule with a parity-twirl argument to show that all two-point cross-correlations between regions A and B vanish for separable, physical states; by Wick's theorem, this then eliminates *all* correlations in the Gaussian case.

Consider a separable fermionic state with decomposition

$$\rho^{AB} = \sum_i p_i \, \rho_i^A \otimes \rho_i^B.$$

A sufficient condition for ρ^{AB} to satisfy the fermionic superselection rule is that each component respects parity:

$$[\rho_i^{AB}, P^{AB}] = 0 \quad \text{with} \quad \rho_i^{AB} := \rho_i^A \otimes \rho_i^B.$$

Equivalently, each factor commutes with the local parity,

$$[\rho_i^A, P^A] = 0, \qquad [\rho_i^B, P^B] = 0 \quad \forall i.$$

Let γ_m^A and γ_n^B be Majorana operators on A and B. Using $P^A \gamma_m^A P^A = -\gamma_m^A$, $(P^A)^2 = I^A$, and the cyclicity of trace, we compute

$$\begin{aligned}
\operatorname{Tr}\big[(\gamma_m^A \gamma_n^B)\, \rho^{AB}\big] &= \sum_i p_i \, \operatorname{Tr}\big[(\gamma_m^A \gamma_n^B)(\rho_i^A \otimes \rho_i^B)\big] \\
&= \sum_i p_i \, \operatorname{Tr}\big[P^A (\gamma_m^A \gamma_n^B)(\rho_i^A \otimes \rho_i^B) P^A\big] \\
&= \sum_i p_i \, \operatorname{Tr}\big[-(\gamma_m^A \gamma_n^B)\, P^A (\rho_i^A \otimes \rho_i^B) P^A\big] \\
&= -\sum_i p_i \, \operatorname{Tr}\big[(\gamma_m^A \gamma_n^B)(\rho_i^A \otimes \rho_i^B)\big] \\
&= -\operatorname{Tr}\big[(\gamma_m^A \gamma_n^B)\, \rho^{AB}\big],
\end{aligned}$$

where in the third line we used $P^A \gamma_m^A P^A = -\gamma_m^A$, and in the fourth line we used that P^A commutes with $\rho_i^A \otimes \rho_i^B$ because $[\rho_i^A, P^A] = 0$. Hence

$$\operatorname{Tr}\big[(\gamma_m^A \gamma_n^B)\, \rho^{AB}\big] = 0.$$

Therefore, in any separable fermionic state obeying the superselection rule, all two-point cross-correlations between A and B vanish. For a fermionic *Gaussian* state ρ_{G}^{AB}, Wick's theorem implies that second moments determine all higher moments; thus the absence of cross-correlations implies the absence of correlations of any order between A and B, and consequently

$$\rho_{\mathrm{G}}^{AB} = \rho_{\mathrm{G}}^A \otimes \rho_{\mathrm{G}}^B. \tag{E.60}$$

This conclusion is specific to Gaussian states: in general, non-Gaussian separable states may exhibit classical correlations beyond second order even in the absence of entanglement. □

Corollary E.7. *A fermionic Gaussian state with CM M_{AB} is separable if and only if $M_{AB} = M_A \oplus M_B$.*

Proof. The forward direction is immediate. For the converse, Theorem 5 implies that all two-point correlations between A and B vanish for separable Gaussian states, so the CM is block-diagonal, $M_{AB} = M_A \oplus M_B$. □

Remark E.8. *This motivates a particularly simple entanglement indicator for fermionic Gaussian states:*

$$E_{\mathrm{cq}}(\rho_G^{AB}) := \tfrac{1}{2}\, I(A;B)_{\rho_G}, \tag{E.61}$$

where $I(A;B) = S(A) + S(B) - S(AB)$ is the quantum mutual information. This quantity (i) vanishes iff the state is separable, (ii) coincides with the entropy of entanglement for pure states, (iii) is efficiently computable from the CM, and (iv) is monotone under Gaussian LOCC. In this sense, E_{cq} cleanly captures the classical–quantum dichotomy *of fermionic Gaussian correlations: every state is either a product state or contains genuine quantum entanglement. We investigate this as an entanglement measure for fermionic Gaussian states in a forthcoming paper.*

E.6 Proof of Theorem 6.6

Before presenting the proof, we contrast our setting with the related bosonic case [122, Theorem 3], noting that while their framework relies on the CM structure specific to bosons, our analysis proceeds differently in the fermionic setting. In the bosonic context, restricting to Gaussian separable states, one can always find a separable Gaussian state whose CM differs from that of the given state by a constant factor. Owing to the commutativity of the two CMs, related techniques can then be applied. The situation for fermionic Gaussian states is markedly different, as we now demonstrate.

Proof of (i): Trace distance

Let $\rho^A = \mathrm{tr}_B(\rho^{AB})$ and $\rho^B = \mathrm{tr}_A(\rho^{AB})$. Since $\rho^A \otimes \rho^B \in \mathrm{SEP}(A:B)$, we immediately obtain

$$\|\rho_G^{AB} - \mathrm{SEP}(A{:}B)\|_1 \;\le\; \|\rho_G^{AB} - \rho_{\mathrm{G}}^A \otimes \rho_{\mathrm{G}}^B\|_1. \tag{E.62}$$

The remaining task is to bound $\|\rho_G^{AB} - \rho_{\mathrm{G}}^A \otimes \rho_{\mathrm{G}}^B\|_1$. By Lemma E.2,

$$\|\rho_G^{AB} - \rho_{\mathrm{G}}^A \otimes \rho_{\mathrm{G}}^B\|_1 \;\leq\; \tfrac{1}{2}\|M_{\rho_{\mathrm{G}}^{AB}} - M_{\rho_{\mathrm{G}}^A \otimes \rho_{\mathrm{G}}^B}\|_1. \tag{E.63}$$

Here, $M_{\rho^A \otimes \rho^B}$ retains the diagonal blocks M_A and M_B of $M_{\rho^{AB}}$ but has zero off-diagonal terms. This corresponds to a separable state with no inter-system correlations. Thus

$$M_{\rho_{\mathrm{G}}^{AB}} - M_{\rho_{\mathrm{G}}^A \otimes \rho_{\mathrm{G}}^B} = \begin{pmatrix} 0 & X \\ -X^T & 0 \end{pmatrix}.$$

The trace norm of the block matrix is determined entirely by the off-diagonal block X. We have

$$\begin{aligned} \tfrac{1}{2}\,\|M_{\rho^{AB}} - M_{\rho^A \otimes \rho^B}\|_1 &= \frac{1}{2}\left\| \begin{pmatrix} 0 & X \\ -X^T & 0 \end{pmatrix} \right\|_1 \\ &= \frac{\|X\|_1 + \|X^T\|_1}{2} \\ &= \|X\|_1. \end{aligned}$$

By Lemma 3 of main text, valid for (k_1, k_2)-extendible states,

$$X^T X \;\leq\; \tfrac{4}{k_1 k_2}\, I, \tag{E.64}$$

implying that every singular value of X or X^T is at most $2/\sqrt{k_1 k_2}$. Since there are at most $\min(n_A, n_B)$ nonzero singular values, we obtain the bound

$$\tfrac{1}{2}\,\|M_{\rho^{AB}} - M_{\rho^A \otimes \rho^B}\|_1 \;\leq\; \frac{2\,\min(n_A, n_B)}{\sqrt{k_1 k_2}}. \tag{E.65}$$

Therefore,

$$\|\rho_G^{AB} - \mathrm{SEP}(A\!:\!B)\|_1 \;\leq\; \frac{2\,\min(n_A, n_B)}{\sqrt{k_1 k_2}} = T, \tag{E.66}$$

which establishes the trace-distance bound.

Continuity bound for mutual information

The mutual information can be expressed as

$$I(A;B)_{\rho^{AB}} = S(\rho^A \otimes \rho^B) - S(\rho^{AB}). \tag{E.67}$$

Let

$$T := \|\rho^A \otimes \rho^B - \rho^{AB}\|_1.$$

The Fannes–Audenaert inequality [11], states that for two states ρ, σ on a d-dimensional Hilbert space with $\frac{1}{2}\|\rho - \sigma\|_1 \leq \epsilon$ and $0 \leq \epsilon \leq 1$,

$$|S(\rho) - S(\sigma)| \;\leq\; \epsilon \log_2 d + h(\epsilon),$$

where $h(\epsilon) = -\epsilon \log_2 \epsilon - (1-\epsilon)\log_2(1-\epsilon)$ is the binary entropy. Applying this with $d = 2^{n_A+n_B}$ and $\epsilon = T/2$ gives

$$|S(\rho^A \otimes \rho^B) - S(\rho^{AB})| \;\leq\; \tfrac{1}{2}(n_A + n_B)T + h(T/2). \tag{E.68}$$

Hence,

$$I(A:B) \;\leq\; \tfrac{1}{2}(n_A + n_B)T + h(T/2). \tag{E.69}$$

To ensure the condition $0 \leq \epsilon \leq 1$, we define

$$T \;=\; \min\!\left(\frac{2\,\min(n_A, n_B)}{\sqrt{k_1 k_2}},\, 2\right) \;=\; \frac{2}{\sqrt{k_1 k_2}}\,\min\!\big(n_A, n_B, \sqrt{k_1 k_2}\big). \tag{E.70}$$

This choice is consistent with the fact that the trace norm distance is at most 2, so we simply set $T = 2$ whenever the bound would exceed this value.

Proof of (ii): Relative entropy of entanglement

The relative entropy of entanglement is defined as

$$E_R(\rho^{AB}) := \inf_{\sigma \in \mathrm{SEP}(A:B)} D(\rho^{AB} \,\|\, \sigma), \tag{E.71}$$

where $D(\rho\|\sigma) = \mathrm{Tr}[\rho(\log\rho - \log\sigma)]$.

Since the infimum is taken over all separable states, it is upper bounded by choosing $\sigma = \rho^A \otimes \rho^B$:

$$E_R(\rho_G^{AB}) \;\leq\; D\big(\rho_G^{AB} \,\big\|\, \rho_{\mathrm{G}}^A \otimes \rho_{\mathrm{G}}^B\big)\,. \tag{E.72}$$

But the right-hand side is exactly the quantum mutual information

$$I(A:B) := S(\rho_A) + S(\rho_B) - S(\rho^{AB}), \tag{E.73}$$

so that

$$E_R(\rho_G^{AB}) \;\leq\; I(A;B). \tag{E.74}$$

Using the continuity bound derived below, we further obtain

$$E_R(\rho_G^{AB}) \;\leq\; \tfrac{1}{2}(n_A + n_B)T + h(T/2). \tag{E.75}$$

Proof of (iii): Squashed entanglement

The squashed entanglement is defined as

$$E_{\mathrm{sq}}(\rho^{AB}) := \tfrac{1}{2} \inf_{\rho^{ABE}} I(A:B|E), \tag{E.76}$$

where the infimum is over extensions ρ^{ABE} of ρ^{AB}, and

$$I(A:B|E) = S(AE) + S(BE) - S(ABE) - S(E) \tag{E.77}$$

is the conditional mutual information.

By monotonicity of conditional mutual information,

$$I(A:B|E) \leq I(A:B). \tag{E.78}$$

Thus

$$E_{\mathrm{sq}}(\rho_G^{AB}) \;\leq\; \tfrac{1}{2} I(A:B). \tag{E.79}$$

Applying the continuity bound below then gives

$$E_{\mathrm{sq}}(\rho_G^{AB}) \;\leq\; \tfrac{1}{4}(n_A + n_B)T + \tfrac{1}{2} h(T/2). \tag{E.80}$$

E.7 Construction of Extendible States

In this section we construct a family of two–mode fermionic Gaussian states. We begin by analyzing their extendibility properties and then use them to derive lower bounds, with particular attention to the case of two–sided extendibility. In this section and the next, we denote a (k_1, k_2)–extendible state by $\rho(k_1, k_2)$ and its associated CM by $M(k_1, k_2)$.

E.7.1 Beam splitter and fermionic pure–loss channel

Consider the two–mode unitary beam splitter U_λ, defined by its action on the annihilation operators $\hat{c}_1, \hat{c}_2$ of modes 1 and 2:

$$U_\lambda^\dagger \hat{c}_1 U_\lambda = \sqrt{\lambda}\, \hat{c}_1 + \sqrt{1-\lambda}\, \hat{c}_2, \tag{E.81}$$

$$U_\lambda^\dagger \hat{c}_2 U_\lambda = \sqrt{\lambda}\, \hat{c}_2 - \sqrt{1-\lambda}\, \hat{c}_1. \tag{E.82}$$

If the second input port is prepared in the vacuum, the first output mode undergoes the fermionic pure–loss channel with transmissivity λ, denoted $\mathcal{E}_\lambda(\cdot)$. For any single–mode input state $\rho = \sum_{m,n=0}^{1} \rho_{m,n} \, |m\rangle\langle n|$, the action is

$$\mathcal{E}_\lambda(|m\rangle\langle n|) = \sum_{\ell=0}^{\min\{m,n\}} \sqrt{\binom{m}{\ell}\binom{n}{\ell}} \, (1-\lambda)^\ell \, \lambda^{\frac{m+n}{2}-\ell} \, |m-\ell\rangle\langle n-\ell| \, .$$

E.7.2 From EPR pairs to (k_1, k_2)–extendible states

We now use the pure–loss channel to construct a (k_1, k_2)–extendible state. Starting with an EPR pair, we apply a beam splitter of parameter λ_1 to subsystem A and another of parameter λ_2 to subsystem B. Tracing through the pure–loss channel action, the EPR state evolves to

$$\begin{aligned} &\tfrac{1}{2}\Big(\big(1 + (1-\lambda_1)(1-\lambda_2)\big) \, |00\rangle\langle 00| + \lambda_2(1-\lambda_1) \, |01\rangle\langle 01| \\ &\quad + \lambda_1(1-\lambda_2) \, |10\rangle\langle 10| + \sqrt{\lambda_1\lambda_2} \, (|00\rangle \langle 11| + |11\rangle \langle 00|) + \lambda_1\lambda_2 \, |11\rangle\langle 11| \Big). \end{aligned}$$

Choosing $\lambda_1 = 1/k_1$ and $\lambda_2 = 1/k_2$ yields the CM

$$M(k_1, k_2) = \begin{pmatrix} 0 & \frac{k_1-1}{k_1} & 0 & \frac{1}{\sqrt{k_1 k_2}} \\ -\frac{k_1-1}{k_1} & 0 & \frac{1}{\sqrt{k_1 k_2}} & 0 \\ 0 & -\frac{1}{\sqrt{k_1 k_2}} & 0 & \frac{k_2-1}{k_2} \\ -\frac{1}{\sqrt{k_1 k_2}} & 0 & -\frac{k_2-1}{k_2} & 0 \end{pmatrix}. \tag{E.83}$$

For $k_1 = k_2 = 1$ this reduces to the original EPR CM, which is not extendible.

E.7.3 Bona fide condition

The matrix $M(k_1, k_2)$ is real and antisymmetric; to be a valid CM it must also satisfy the fermionic bona fide condition, i.e. the spectrum of $iM(k_1, k_2)$ lies within $[-1, 1]$.

- If $k_1 = 1$ (right–extendibility only), the eigenvalues are $\{\pm 1, \pm 1/k_2\}$, all within $[-1, 1]$. - If $k_2 = 1$ (left–extendibility only), the spectrum is $\{\pm 1, \pm 1/k_1\}$, again valid. - The

two–sided case requires a more detailed bound. Define

$$\widetilde{M}(k_1,k_2)=\begin{pmatrix} 0 & (k_1-1)\sqrt{\frac{k_2}{k_1}} & 0 & 1 \\ -(k_1-1)\sqrt{\frac{k_2}{k_1}} & 0 & 1 & 0 \\ 0 & -1 & 0 & (k_2-1)\sqrt{\frac{k_1}{k_2}} \\ -1 & 0 & -(k_2-1)\sqrt{\frac{k_1}{k_2}} & 0 \end{pmatrix}, \tag{E.84}$$

so that $M(k_1,k_2)=(1/\sqrt{k_1k_2})\,\widetilde{M}(k_1,k_2)$. Writing $a=(k_1-1)\sqrt{k_2/k_1}$ and $b=(k_2-1)\sqrt{k_1/k_2}$, the eigenvalues of $i\widetilde{M}(k_1,k_2)$ are $\pm ir_1,\pm ir_2$ with

$$r_{1,2}^2=\tfrac{a^2+b^2}{2}+1\pm\tfrac{1}{2}|a-b|\sqrt{a^2+2ab+b^2+4}. \tag{E.85}$$

Thus $r_1^2+r_2^2=a^2+b^2+2$, implying $\max\{|r_1|,|r_2|\}\geq 1$. This means that the largest singular value of $M(k_1,k_2)$ is lower bounded by $1/\sqrt{k_1k_2}$, so the moduli of its eigenvalues lie in

$$(-\infty,-1/\sqrt{k_1k_2}]\cup[1/\sqrt{k_1k_2},\infty).$$

On the other hand, this lower bound is achieved when $a=b=0$, i.e. in the non–extendible case corresponding to the original EPR pair with eigenvalues $\pm i,\pm i$. Hence the operator norm is exactly $1/\sqrt{k_1k_2}$, and the eigenvalues of $M(k_1,k_2)$ lie in the interval

$$\big[-i/\sqrt{k_1k_2},\,i/\sqrt{k_1k_2}\big],$$

thereby satisfying the fermionic bona fide condition.

E.7.4 Verification of (k_1,k_2)–extendibility

To show (k_1,k_2)–extendibility explicitly, we use Theorem 2. Here

$$M_A=\begin{pmatrix} 0 & \frac{k_1-1}{k_1} \\ -\frac{k_1-1}{k_1} & 0 \end{pmatrix},\quad M_B=\begin{pmatrix} 0 & \frac{k_2-1}{k_2} \\ -\frac{k_2-1}{k_2} & 0 \end{pmatrix},$$
$$X=\begin{pmatrix} 0 & \frac{1}{\sqrt{k_1k_2}} \\ \frac{1}{\sqrt{k_1k_2}} & 0 \end{pmatrix}.$$

The theorem requires the existence of antisymmetric matrices Y, Z with $I + iY \geq 0$ and $I + iZ \geq 0$ such that

$$\begin{pmatrix} M_A - Z & 0 & 0 & 0 \\ 0 & M_A + (k_1 - 1)Z & \sqrt{k_1 k_2} X & 0 \\ 0 & -\sqrt{k_1 k_2} X^T & M_B + (k_2 - 1)Y & 0 \\ 0 & 0 & 0 & M_B - Y \end{pmatrix}$$

is a valid CM.

- For $k_1 = k_2 = 1$, choosing $Y = Z = 0$ suffices. - For all other cases, one may choose

$$Z = \begin{pmatrix} 0 & -1/k_1 \\ 1/k_1 & 0 \end{pmatrix}, \quad Y = \begin{pmatrix} 0 & -1/k_2 \\ 1/k_2 & 0 \end{pmatrix}.$$

With these choices the resulting CM corresponds to an EPR pair, and is therefore legitimate. This confirms (k_1, k_2)–extendibility.

E.8 Lower Bounds on the Distance to Fermionic Separable States

Here we employ the (k_1, k_2)-extendible state constructed in the previous section to derive lower bounds on the distance to fermionic separable states. We first present a bound that is tight. We then show that adapting the bosonic strategy of [122] does not yield a tight lower bound in the fermionic setting; instead, the resulting estimate is weaker by an additional factor of two.

E.8.1 Tight Lower Bound from Separable States

Theorem E.9. *Let ρ be a fermionic Gaussian state with CM define above $M(k_1, k_2)$. Then, the trace distance between ρ and the set of (possibly non-Gaussian) separable fermionic states satisfies the lower bound*

$$\|\rho - \mathrm{SEP}(A : B)\|_1 \geq \frac{1}{\sqrt{k_1 k_2}}.$$

Proof. Applying Lemma E.1, we obtain

$$\|M_\rho - M_\sigma\|_\infty \leq \|\rho - \sigma\|_1.$$

Minimizing both sides independently over all separable states σ, we get

$$\inf_{M_\sigma : \sigma \in \mathrm{SEP}} \|M_\rho - M_\sigma\|_\infty \leq \|\rho - \mathrm{SEP}\|_1.$$

Note that the optimal σ minimizing each side may differ. Our goal is to obtain a lower bound on the left-hand side. In fact, we show that the infimum can be computed exactly.

From Theorm 5, any separable (possibly non-Gaussian) fermionic state σ has a correlation matrix of the form

$$M_\sigma = \begin{pmatrix} 0 & a & 0 & 0 \\ -a & 0 & 0 & 0 \\ 0 & 0 & 0 & b \\ 0 & 0 & -b & 0 \end{pmatrix},$$

with $|a|, |b| \leq 1$. Then,

$$\inf_{M_\sigma : \sigma \in \mathrm{SEP}} \|M_\rho - M_\sigma\|_\infty = \inf_{|a|\leq 1, |b| \leq 1} \left\| \begin{pmatrix} 0 & \frac{k_1 - 1}{k_1} - a & 0 & \frac{1}{\sqrt{k_1 k_2}} \\ -\left(\frac{k_1 - 1}{k_1} - a\right) & 0 & \frac{1}{\sqrt{k_1 k_2}} & 0 \\ 0 & -\frac{1}{\sqrt{k_1 k_2}} & 0 & \frac{k_2 - 1}{k_2} - b \\ -\frac{1}{\sqrt{k_1 k_2}} & 0 & -\left(\frac{k_2 - 1}{k_2} - b\right) & 0 \end{pmatrix} \right\|_\infty.$$

Let us define

$$\bar{a} := \left(\frac{k_1 - 1}{k_1} - a\right)\sqrt{k_1 k_2}, \quad \bar{b} := \left(\frac{k_2 - 1}{k_2} - b\right)\sqrt{k_1 k_2}.$$

Translating the constraints $|a| \leq 1$ and $|b| \leq 1$ to conditions on $\bar{a}, \bar{b}$, the problem becomes

$$\frac{1}{\sqrt{k_1 k_2}} \inf_{\substack{-\sqrt{\frac{k_2}{k_1}} \leq \bar{a} \leq 2\sqrt{k_1 k_2} - \sqrt{\frac{k_2}{k_1}} \\ -\sqrt{\frac{k_1}{k_2}} \leq \bar{b} \leq 2\sqrt{k_1 k_2} - \sqrt{\frac{k_1}{k_2}}}} \left\| \begin{pmatrix} 0 & \bar{a} & 0 & 1 \\ -\bar{a} & 0 & 1 & 0 \\ 0 & -1 & 0 & \bar{b} \\ -1 & 0 & -\bar{b} & 0 \end{pmatrix} \right\|_\infty = \frac{1}{\sqrt{k_1 k_2}}. \tag{E.86}$$

The eigenvalues λ_1, λ_2 of a real antisymmetric 4×4 matrix come in purely imaginary conjugate pairs $\pm i r_1, \pm i r_2$, with squared moduli given by

$$\lambda_{1,2}^2 = -\frac{\bar{a}^2 + \bar{b}^2}{2} \pm \frac{1}{2}(\bar{a} - \bar{b})\sqrt{\bar{a}^2 + 2\bar{a}\bar{b} + \bar{b}^2 + 4} - 1. \tag{E.87}$$

The operator norm is the largest singular value, i.e., $\max\{|r_1|, |r_2|\}$, and since

$$r_1^2 + r_2^2 = \bar{a}^2 + \bar{b}^2 + 2,$$

we conclude that

$$\max\{|r_1|, |r_2|\} \geq 1.$$

This lower bound is achieved when $\bar{a} = \bar{b} = 0$, yielding eigenvalues $\pm i, \pm i$. Hence, the operator norm is exactly 1, and the infimum in Eq. (E.86) is achieved, completing the proof. □

Proposition E.10. *The trace-distance bound in Theorem E.9 is tight for two-mode fermionic Gaussian states. Specifically, there exist (k_1, k_2)-extendible Gaussian state ρ, such that:*

$$\frac{1}{\sqrt{k_1 k_2}} \leq \|\rho - \mathrm{SEP}(A:B)\|_1 \leq \frac{2}{\sqrt{k_1 k_2}}.$$

Proof. The lower bound follows from Theorem E.9, and the upper bound from Theorem 6. □

E.8.2 Adapting the Bosonic Strategy

To estimate $\|\rho(k_1, k_2) - \mathrm{SEP}\|_1$, we first recall that for any traceless operator X it holds that $\|X\|_1 \geq 2\|X\|_\infty$. This directly yields the lower bound

$$\|\rho(k_1, k_2) - \mathrm{SEP}\|_1 \geq 2 \|\rho(k_1, k_2) - \mathrm{SEP}\|_\infty . \tag{E.88}$$

From Lemma E.5 we know that, for every bipartite pure state $|\Psi\rangle$ with maximal Schmidt coefficient $\lambda_1(\Psi)$, the inequality $\langle\Psi|\sigma|\Psi\rangle \leq \lambda_1(\Psi)$ holds for all separable states σ. Hence,

$$\|\rho(k_1, k_2) - \mathrm{SEP}\|_\infty = \inf_{\sigma \in \mathrm{SEP}} \|\rho(k_1, k_2) - \sigma\|_\infty , \tag{E.89}$$

$$= \inf_{\sigma \in \mathrm{SEP}} \sup_{|\Psi\rangle} |\langle\Psi| (\rho(k_1, k_2) - \sigma) |\Psi\rangle| , \tag{E.90}$$

$$\geq \sup_{|\Psi\rangle} (\langle\Psi| \rho(k_1, k_2) |\Psi\rangle - \lambda_1(\Psi)) . \tag{E.91}$$

We now specialize to the case where $|\Psi\rangle$ is chosen as the EPR pair $\frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$, which satisfies $\lambda_1(\Psi) = \frac{1}{2}$. In this case, $\langle\Psi|\rho(k_1, k_2)|\Psi\rangle = \mathrm{tr}(\rho(k_1, k_2)\rho(1, 1))$.

For two fermionic Gaussian states $\rho(k_1, k_2)$ and $\rho(1,1)$ with CMs $M(k_1, k_2)$ and $M(1,1)$, respectively, one has the identity

$$\left|\operatorname{tr}(\rho(k_1,k_2)\rho(1,1))\right| = \sqrt{\det\left(\tfrac{1}{2}\left(M(k_1,k_2)M(1,1) - I\right)\right)}. \tag{E.92}$$

Applying this relation, we obtain

$$\sqrt{\det\left(\tfrac{1}{2}\left(M(k_1,k_2)M(1,1) - I\right)\right)} = \frac{1}{4}\left(\frac{2k_1k_2 - k_1 - k_2 + 2}{k_1k_2} + \frac{2}{\sqrt{k_1k_2}}\right) - \frac{1}{2}. \tag{E.93}$$

It is straightforward to verify that this bound is always weaker than the one derived in Proposition E.9. For the symmetric case $k_1 = k_2 = k$, the expression simplifies to

$$\sqrt{\det\left(\tfrac{1}{2}\left(M(k,k)M(1,1) - I\right)\right)} = \frac{1}{2}\left(1 + \frac{1}{k^2}\right), \tag{E.94}$$

which leads to

$$\|\rho(k) - \mathrm{SEP}\|_1 \geq \frac{1}{2k^2}. \tag{E.95}$$

Recall that the bound from Theorem E.9 in the same setting is $\frac{1}{k}$. Therefore,

$$\frac{1}{k} \geq \frac{1}{2k^2}, \quad \text{for all } k \geq 1. \tag{E.96}$$

We emphasize that this inequality reflects more than a difference in constant prefactors: it highlights a change in scaling behavior, with the first bound decaying linearly in k, while the second decays quadratically.

E.9 Numerical illustration: disordered Kitaev chain

To complement the general covariance-matrix characterization developed in the main text, we include a simple many-body example based on a disordered quadratic fermionic model. The purpose of this appendix is not to analyze the phase diagram of the model, but rather to illustrate concretely how our extendibility criterion can be applied to reduced states arising from a standard many-body system.

Specifically, we ask the following question: given the ground state of a disordered free-fermion Hamiltonian, how often is the reduced state on two neighboring sites (k_1, k_2)-extendible? Since extendibility provides a natural way to probe how local correlations

are modified, we apply this to understand how localized the correlations are as disorder increases. For the example considered below, we find that the probability of extendibility increases with disorder strength, consistent with the intuition that stronger disorder suppresses inter-site correlations and tends to localize them.

As a concrete testbed, we consider the standard disordered Kitaev chain, namely a one-dimensional spinless p-wave superconductor. At the symmetric point $t = \Delta = 1$, the Hamiltonian on a ring of length N is

$$H \;=\; \sum_{j=1}^{N} \Big[-(c_j^\dagger c_{j+1} + c_{j+1}^\dagger c_j) \;+\; (c_j c_{j+1} + c_{j+1}^\dagger c_j^\dagger) \Big] \;-\; \sum_{j=1}^{N} \mu_j \Big(c_j^\dagger c_j - \tfrac{1}{2} \Big), \tag{E.97}$$

with periodic boundary conditions ($c_{N+1} \equiv c_1$). The onsite chemical potentials are sampled independently from the uniform distribution

$$\mu_j \sim \mathrm{Unif}[-W, W],$$

where $W \geq 0$ denotes the disorder strength.

Because the model is quadratic, its many-body ground state is a fermionic Gaussian state and is therefore completely determined by its Majorana covariance matrix. Writing the Majorana operators as $\gamma = (\gamma_1, \ldots, \gamma_{2N})^{\mathsf{T}}$, any quadratic Hamiltonian can be expressed as

$$H \;=\; \frac{i}{4}\, \gamma^{\mathsf{T}} A\, \gamma, \tag{E.98}$$

where A is a real antisymmetric $2N \times 2N$ matrix. For a state ρ, the corresponding Majorana covariance matrix is defined by

$$\Gamma_{pq} \;:=\; \frac{i}{2}\, \mathrm{Tr}\big(\rho\, [\gamma_p, \gamma_q]\big) \;=\; \mathrm{Tr}\big(\rho\, i\gamma_p \gamma_q\big) \qquad (p \neq q), \qquad \Gamma_{pp} = 0. \tag{E.99}$$

For a pure Gaussian ground state, Γ is real antisymmetric and satisfies $\Gamma^2 = -I$.

In the numerics, we compute the ground state of the Hamiltonian Γ from the quadratic generator $K := iA$ via

$$\Gamma \;=\; i\, \mathrm{sgn}(K), \tag{E.100}$$

where $\mathrm{sgn}(K)$ is obtained by diagonalizing K and replacing each eigenvalue by its sign.

For each value of the disorder strength W, we sample N_{samp} independent realizations and define

$$p(W) \;:=\; \Pr_{\{\mu_j\}} \Big[\rho_G^{(12)} \text{ is } (k_1, k_2)\text{-extendible} \Big] \;\approx\; \frac{N_{\mathrm{ext}}(W)}{N_{\mathrm{samp}}}, \tag{E.101}$$

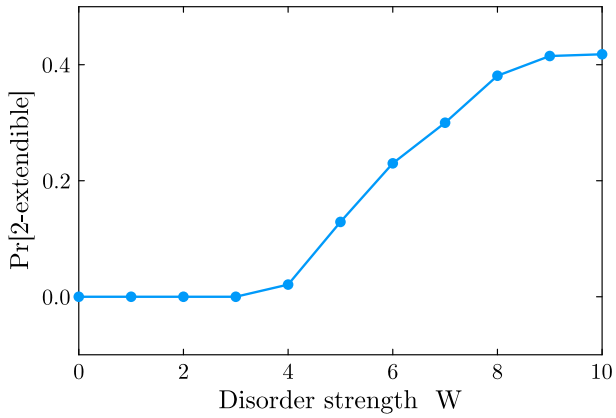


Figure E.1: **Disordered Kitaev chain: probability of two-site (k_1, k_2)-extendibility versus disorder.** For the disordered Kitaev chain (E.97) with periodic boundary conditions and i.i.d. chemical-potential disorder $\mu_j \sim \mathrm{Unif}[-W, W]$, we compute the Gaussian ground-state covariance matrix, reduce it to two adjacent sites, and test two-sided (k_1, k_2)-extendibility of the resulting two-site Gaussian state using the covariance-matrix SDP characterization from the main text. The plotted quantity is the empirical probability $p(W)$ over disorder realizations. The rise of $p(W)$ with W indicates that stronger disorder makes the local reduced state more likely to satisfy the extendibility constraints. Parameters: $N = 10$, $(k_1, k_2) = (1, 2)$, and $N_{\mathrm{samp}} = 1000$.

where $N_{\mathrm{ext}}(W)$ is the number of sampled disorder realizations, out of N_{samp}, for which the reduced two-site state is (k_1, k_2)-extendible.

The figure shows the empirical probability $p(W)$ as a function of W, for fixed (k_1, k_2) and fixed global system size N. We stress that extendibility is tested only for the reduced two-site state; the role of the full system size N is merely to define the parent many-body Gaussian ground state from which this local reduction is obtained. Since the quantity under consideration is local, one expects the curve to become essentially insensitive to N once N is sufficiently large.

A notable feature of the plot is the existence of a low-disorder regime in which $p(W) = 0$, meaning that none of the sampled realizations yield a (k_1, k_2)-extendible reduced state. As W increases, a nonzero extendible fraction emerges, indicating a crossover—or, at the numerical level, an apparent transition—from a regime in which the local correlations are too strong to satisfy the extendibility constraints to one in which extendibility occurs with positive probability.

This behavior is consistent with the interpretation of extendibility as a constraint on

the strength and structure of bipartite correlations. Increasing disorder suppresses the effective cross-correlations across a fixed local cut, thereby making the reduced two-site state progressively more likely to satisfy the (k_1, k_2)-extendibility conditions. In this way, the plot provides a concrete illustration of how the covariance-matrix extendibility criterion developed in the main text can be used to study the local correlations in disordered fermionic Gaussian systems.