

Ergodicity breaking and stabilisation of quantum order

Jared Jeyaretnam

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
of
University College London.

Department of Physics and Astronomy
University College London

19th December 2023

I, Jared Jeyaretnam, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the work.

Abstract

At zero temperature, symmetry-protected topological (SPT) order can encode quantum information in an edge strong zero mode, robust to perturbations respecting some symmetry. On the other hand, phenomena like many-body localisation (MBL) and quantum scarring can arrest the approach to thermal equilibrium, contrary to the ergodic dynamics expected of generic quantum systems. This raises the possibility that by combining SPT order with such ergodicity breaking phenomena, one might be able to construct a quantum memory that is robust at finite temperature. However, MBL is threatened by thermal avalanches due to rare low-disorder grains, while the survival of quantum scars in the thermodynamic limit is still uncertain beyond certain finely-tuned models. To this end, we study the interaction between ergodicity breaking and novel quantum order in spin chains. First, we look at an interacting spin-1/2 chain hosting SPT order and protected edge modes, identify quantum scars responsible for long-lived coherence in the bulk of the system, and show that these scars exhibit signatures of SPT order even at finite energy density. Secondly, we study a topological transition between two MBL phases with different topological order. Through a renormalisation group approach, we identify many-body resonances between the localised states; close to the transition these proliferate, link together into an extensive network, and cause delocalisation. Additionally, we characterise the SPT strong zero mode in the ordered phase. This has important implications for the stability of MBL and transitions between MBL phases with different topological orders. Finally, we investigate a model known for its use in a proof of MBL, finding that the localised phase is destabilised by proximity to an SPT critical point even though the system itself has no symmetries. This thesis therefore provides important insight into the stability of quantum order at finite temperatures.

Impact Statement

Emerging quantum technologies, in particular quantum computing, have widespread applications including in physical simulation, drug discovery, quantum chemistry, novel materials, and cryptography, and companies such as IBM, Google, and Microsoft are investing billions into the race to develop viable quantum computers. A major challenge is the preservation of quantum information over long timescales: the state of the art in quantum memories is on the order of single milliseconds, which severely limits the length and accuracy of computations. In this thesis, we consider two forms of ergodicity breaking which enable systems to avoid thermal equilibrium: many-body localisation (MBL), in which quenched disorder freezes the dynamics, and quantum many-body scarring, where atypical eigenstates cause long-lived oscillations from certain initial states. We show that these stabilise a novel form of quantum order known as symmetry-protected topological (SPT) order, and in the case of MBL, probe the stability of the MBL phase itself near topological phase transitions. SPT order can be used to encode quantum information in a robust manner, for example in an edge mode, while disorder is crucial to the operation of certain classes of quantum computer, such as superconducting transmon qubit arrays. For this reason, this thesis provides important insight into the stabilisation of quantum information, and the development of quantum technologies, even though these applications are not the direct focus of the research. On the academic side, ergodicity breaking is important to the field of quantum many-body physics, as it challenges our assumptions about conventional statistical mechanics and enables us to observe novel out-of-equilibrium phenomena. Chapters 3 and 4 respectively are based upon research which has been published in two separate papers in Physical Review B [1, 2]; the former was the first to demonstrate SPT order in quantum scars without fine-tuning, while the latter provides new insight into an ongoing debate about the stability of MBL

in the thermodynamic limit. The results contained within this thesis have also been disseminated through conferences, workshops, and talks, reaching not just the condensed matter community but also the quantum information community. In Chapter 3, we also develop an extension to a real-space renormalisation group algorithm, with general applicability to disordered spin chains.

UCL Research Paper Declarations

Form 1

1. **For a research manuscript that has already been published** (if not yet published, please skip to section 2):

- (a) **What is the title of the manuscript?**

Quantum scars and bulk coherence in a symmetry-protected topological phase

- (b) **Please include a link to or doi for the work:**

[10.1103/PhysRevB.104.014424](https://doi.org/10.1103/PhysRevB.104.014424)

- (c) **Where was the work published?**

Physical Review B

- (d) **Who published the work?**

American Physical Society (APS)

- (e) **When was the work published?**

21 July 2021

- (f) **List the manuscript's authors in the order they appear on the publication:**

Jared Jeyaretnam, Jonas Richter, and Arijeet Pal

- (g) **Was the work peer reviewed?**

Yes

- (h) **Have you retained the copyright?**

No

- (i) **Was an earlier form of the manuscript uploaded to a preprint server (e.g. medRxiv)? If 'Yes', please give a link or doi**

Yes: [arXiv:2103.15880](https://arxiv.org/abs/2103.15880)

If 'No', please seek permission from the relevant publisher and check the box next to the below statement:

☐ *I acknowledge permission of the publisher named under 1d to include in this thesis portions of the publication named as included in 1c.*

2. For a research manuscript prepared for publication but that has not yet been published (if already published, please skip to section 3):

(a) **What is the current title of the manuscript?**

(b) **Has the manuscript been uploaded to a preprint server?**

If 'Yes', please give a link or doi:

(c) **Where is the work intended to be published?**

(d) **List the manuscript's authors in the intended authorship order:**

(e) **Stage of publication:**

3. For multi-authored work, please give a statement of contribution covering all authors (if single-author, please skip to section 4):

Jared Jeyaretnam and Arijeet Pal initially conceptualised the work. Jared Jeyaretnam discovered the scars and performed the eigenstate analysis. Jonas Richter characterised the dynamics, finding signatures of bulk coherence. Jared Jeyaretnam and Jonas Richter contributed equally to writing the majority of the manuscript. All authors contributed equally to understanding the results.

4. In which chapter(s) of your thesis can this material be found?

Chapter 3

e-Signatures confirming that the information above is accurate (this form should be co-signed by the supervisor/ senior author unless this is not appropriate, e.g. if the paper was a single-author work):

Candidate: Jared Jeyaretnam

Date: 19th December 2023

Supervisor/Senior Author signature (where appropriate): Arijeet Pal

Date: 29th September 2023

Form 2

1. **For a research manuscript that has already been published** (if not yet published, please skip to section 2):

- (a) **What is the title of the manuscript?**

Renormalization view on resonance proliferation between many-body localized phases

- (b) **Please include a link to or doi for the work:**

[10.1103/PhysRevB.108.094205](https://doi.org/10.1103/PhysRevB.108.094205)

- (c) **Where was the work published?**

Physical Review B

- (d) **Who published the work?**

American Physical Society (APS)

- (e) **When was the work published?**

11 September 2023

- (f) **List the manuscript's authors in the order they appear on the publication:**

Jared Jeyaretnam, Christopher J. Turner, Arijeet Pal

- (g) **Was the work peer reviewed?**

Yes

- (h) **Have you retained the copyright?**

The work is published under the [CC BY 4.0 licence](https://creativecommons.org/licenses/by/4.0/).

- (i) **Was an earlier form of the manuscript uploaded to a preprint server (e.g. medRxiv)? If 'Yes', please give a link or doi**

Yes: [arXiv:2301.08738](https://arxiv.org/abs/2301.08738)

If 'No', please seek permission from the relevant publisher and check the box next to the below statement:

☐ *I acknowledge permission of the publisher named under 1d to include in this thesis portions of the publication named as included in 1c.*

2. **For a research manuscript prepared for publication but that has not yet been published** (if already published, please skip to section 3):

- (a) **What is the current title of the manuscript?**

(b) **Has the manuscript been uploaded to a preprint server?**

If 'Yes', please please give a link or doi:

(c) **Where is the work intended to be published?**

(d) **List the manuscript's authors in the intended authorship order:**

(e) **Stage of publication:**

3. For multi-authored work, please give a statement of contribution covering all authors (if single-author, please skip to section 4):

Jared Jeyaretnam and Arijeet Pal initially conceptualised the work. All three authors contributed equally to the development of techniques and the understanding of the results. Jared Jeyaretnam performed all the calculations and largely wrote the manuscript.

4. In which chapter(s) of your thesis can this material be found?

Chapter [4](#)

e-Signatures confirming that the information above is accurate (this form should be co-signed by the supervisor/ senior author unless this is not appropriate, e.g. if the paper was a single-author work):

Candidate: Jared Jeyaretnam

Date: 19th December 2023

Supervisor/Senior Author signature (where appropriate): Arijeet Pal

Date: 29th September 2023

Acknowledgements

Firstly, I would like to thank my supervisor Arijeet Pal for his excellent guidance, and in particular for suggesting that I investigate quantum scars when I first started, which proved to be a fascinating field of research. Under his mentorship, I've developed from a mere student into a physics researcher. I am also grateful to my collaborators Jonas Richter and Chris Turner for many stimulating discussions, and for helping me grow as a researcher – in particular, Jonas for working with me on my first ever publication, and Chris for always having an answer to my questions. I'd also like to thank all the members of my research group: for the countless great chats – both about physics and otherwise – and for our lunchtime trips to all the best food stands around UCL. I have a lot of fond memories from our office, and I'll miss it.

Of course, I could not have made it through this PhD without the support of so many excellent people in my life. I cannot thank my wonderful girlfriend Alessandra enough for her constant support, and for always putting up with my grumbling about bugs in my code! Her incredible cooking skills alone have kept my spirits up even through difficult research projects, and travelling with her has broadened my horizons. Thank you also to the UCL Climbing and Mountaineering Club, for giving me a fantastic hobby and keeping me in shape at the same time, and especially for being my social lifeline during lockdown. And thank you to my friends, from all of my walks of life, for being there for me.

Finally, my parents: from a very young age, they always encouraged me to be inquisitive and indulged my curiosities, from experiments in the kitchen with Mum, to getting me a *New Scientist* subscription. I still fondly remember going to the Science Museum with my Dad to watch the Curiosity rover land on Mars. Thank you for raising me to be a scientist, for showering me with love, and for always being just a phone call away.

Contents

Abstract	iii
Impact Statement	iv
UCL Research Paper Declaration Forms	vi
Contents	x
List of Figures	xiv
1 Introduction	1
2 Background	5
2.1 Thermalisation and Quantum Chaos	5
2.1.1 The Eigenstate Thermalisation Hypothesis	6
2.1.2 Level Statistics	8
2.1.3 The Quantum Information Perspective	9
2.2 Avoiding Thermalisation: Many Body Localisation	11
2.2.1 Anderson localisation	12
2.2.2 The interacting many-body case	14
2.2.3 Consequences of many-body localisation	18
2.2.4 Breakdown of many-body localisation	21
2.3 Quantum Scars	27
2.3.1 Persistent oscillations in chaotic systems	29
2.3.2 Theories of quantum scarring	32
2.3.3 Connection to single-particle quantum scars	35
2.4 Topological Order	38
2.4.1 Classification of topological order	38
2.4.2 Edge modes, entanglement spectra, and string order	40

2.4.3	Robust topological order at finite temperature	46
3	Quantum scars in an SPT phase	49
3.1	Introduction	49
3.2	Lack of SPT order in AKLT scars	50
3.3	Quantum Scars in the ZXZ Model	52
3.3.1	Model	55
3.3.2	Methods	58
3.3.3	Eigenstate Analysis	59
3.3.4	Level statistics and onset of chaos	70
3.3.5	Non-equilibrium dynamics	71
3.3.6	Discussion	77
3.4	Supplementary Material	80
3.4.1	Exact eigenstates of the clean ZXZ model	80
3.4.2	Matrix Product State representation	81
3.4.3	Non-zero momentum excitations	84
4	Resonance proliferation between MBL phases	86
4.1	Introduction	86
4.2	Model and RSRG-X	89
4.3	Results	95
4.3.1	Spin-glass order	95
4.3.2	Energy variance	97
4.3.3	Many-body resonances	99
4.3.4	Resonant Clusters	102
4.3.5	Scaling of resonance behaviour with system size	104
4.4	Discussion	105
4.5	Supplementary Material	107
4.5.1	RSRG-X Decimation Rules	107
4.5.2	Clifford RSRG-X	108
4.5.3	Distribution of maximal resonant cluster sizes	110
4.5.4	Higher resonance threshold	111

5	Disordered Ising model with local fields	113
5.1	Introduction	113
5.2	Model	114
5.3	RSRG-X on Imbrie's Model	116
5.3.1	Decimation Rules	116
5.3.2	Results	118
5.4	Fermion Perturbation Theory	121
5.4.1	Method	122
5.4.2	Results	124
5.5	Phase Diagram and Discussion	127
6	Conclusions	129
	Bibliography	135

List of Figures

2.1	From Ref. [70]: Local thermalisation results from global unitary evolution, due to the spreading of entanglement.	10
2.2	From Ref. [90]: Spin-spin correlations showing exponential decay in the localised phase.	15
2.3	From Ref. [97]: Interacting cold atoms trapped in a 2D optical lattice display signs of many-body localisation when subjected to quenched disorder, retaining memory of initial conditions	16
2.4	From Ref. [132]: A thermal avalanche, resulting from a rare low-disorder region, may destabilise the entire MBL phase.	23
2.5	From Ref. [132]: A thermalising bath in a 1D Bose-Hubbard model, experimentally realised in an optical lattice, destabilises the MBL phase in a thermal avalanche.	25
2.6	From Ref. [9]: Unexpected persistent oscillations between two charge density wave states in a Rydberg atom simulation	28
2.7	From Ref. [18]: Numerical simulations of the PXP model, which captures the key features of the Rydberg blockade regime	28
2.8	From Ref. [19]: Entanglement entropies and overlaps with the charge density wave state $ \mathbb{Z}_2\rangle$, showing the band of states responsible for the persistent oscillations in Fig. 2.7	30
2.9	From Ref. [153]: A superconducting processor used to experimentally realise quantum scars.	31
2.10	From Ref. [153]: Oscillations in a superconducting qubit model of quantum scarring	33
2.11	From Ref. [20]: Chaotic and periodic trajectories in the Bunimovich stadium, and their relation to quantum scars.	36

2.12	From Ref. [187]: Quasienergy spectroscopy allows the identification of Majorana edge modes, showing that their splitting Δ shrinks exponentially with system size.	44
2.13	From Ref. [187]: Reconstructing the edge modes in an Ising chain.	44
3.1	Ground states of the AKLT model	50
3.2	Cartoons introducing the ZXZ Hamiltonian and main results of this section	54
3.3	Overlaps of ZXZ eigenstates with $ \psi_{\bar{1}/0}\rangle$, showing an atypical band of states with equal energy spacing; and the Loschmidt echo of $ \psi_{\bar{1}/0}\rangle$, showing oscillations at this frequency	63
3.4	Half-system bipartite entanglement entropy of the ZXZ eigenstates, and sub-volume scaling of scar state entanglement entropies	65
3.5	Various plots of cluster expectation values $\langle K_\ell \rangle$, showing atypicality of the scar states	67
3.6	Ordered entanglement energy spectra φ_α of selected exemplary states in the ZXZ model.	69
3.7	Level statistics for the ZXZ model, showing approach to quantum chaos as $L \rightarrow \infty$ for the chosen parameters.	70
3.8	Infinite-temperature autocorrelation functions at the edge and in the bulk for Z_ℓ and K_ℓ in the ZXZ model	72
3.9	Quench dynamics of cluster expectation values for three exemplary cluster-basis states, with and without $\lambda = 1$	73
3.10	Dynamics of the even- and odd-sublattice cluster expectation values, for the initial state $ \psi_2\rangle$ (c.f. Fig. 3.9)	74
3.11	Overlap of the states $ \psi_1\rangle$ - $ \psi_3\rangle$ (c.f. Fig. 3.9) with the eigenstates $ n\rangle$ of $\mathcal{H}$	75
3.12	Time- and space-averaged correlation function and Loschmidt echo in the $\lambda \neq 1$ model, for states in the cluster basis	76
3.13	Matrix product state representation of a tower state $ T_j^\pm(k)\rangle$	82
3.14	Overlaps of the states $ T_j^\pm(k)\rangle$, constructed with a generalised excitation operator, and $ S_j^\pm(k)\rangle$ in analogy with the inset of Fig. 3.3	84

4.1	Representative RSRG-X tree tensor networks and respective l-bits $\{\tau_k\}$, for the paramagnetic, spin glass, and critical phases	92
4.2	Energy scale of the final l-bit ΔE_f , as a function of δ	96
4.3	Disorder-averaged energy variance $\langle \delta H^2 \rangle / \langle H^2 \rangle$ of RSRG-X leaf states across the δ - g plane	98
4.4	Probability density of $\ln(\mathcal{G}_\alpha)$ for various points in the δ - g plane . .	100
4.5	Density of resonances, and probability that any one l-bit is flipped by at least one resonant term, across the δ - g plane	101
4.6	Probability distribution of resonant cluster sizes s_{res} , for various points in the δ - g plane	102
4.7	Disorder-averaged maximum resonant cluster size $\langle \max\{s_{\text{res}}\} \rangle$, across the δ - g plane	103
4.8	Scaling analysis of $\langle \max\{s_{\text{res}}\} \rangle$ against system size L , and the ergodic-MBL phase boundaries derived from this analysis	104
4.9	Distribution of maximal resonant l-bit cluster sizes, $\max\{s_{\text{res}}\}$, across $\delta - g$ plane	110
4.10	Fig. 4.6, but only considering very strong resonances with $\mathcal{G} > 1$. .	112
4.11	Fig. 4.9, but only considering very strong resonances with $\mathcal{G} > 1$. .	112
5.1	From Ref. [240]: Schematic ground-state RG phase diagram of H_{Imbrie} .	115
5.2	Comparisons between RSRG-X and exact diagonalisation in the TFIM with longitudinal fields.	118
5.3	Histograms of RSRG-X errors in TFIM with longitudinal fields, for various model parameter choices.	120
5.4	RSRG-X errors as a function of the size of small parameters in TFIM with longitudinal fields.	120
5.5	Critical field λ^* as a function of δ , for various system sizes L	125
5.6	Critical field λ^* against the inverse system size $1/L$, for selected values of δ	125
5.7	Critical field λ^* in the $L \rightarrow \infty$ limit, as a function of δ , giving a rough phase diagram.	126
5.8	Full phase diagram for the disordered TFIM with longitudinal fields.	128

Chapter 1

Introduction

Ordinary matter thermalises – that is to say, given any initial state, such as a cup of hot coffee sitting in a cold room or a drop of dye placed into a tank of water, the passing of time will eventually result in a state at thermal equilibrium, described only by global conservation laws and with no memory of detailed initial conditions. The coffee will reach room temperature, and the dye will spread throughout the water. This equilibrium state will be described only by macroscopic conserved quantities such as energy and particle number, while the entropy will tend towards the maximum permitted by said constraints. Despite this empirical fact, it is still not entirely clear how and when this process proceeds for quantum systems: indeed we have discovered mechanisms that prevent the approach to equilibrium entirely, or even more bizarrely cause a return to the initial conditions after apparent thermalisation.

For a generic quantum system, our intuition carries over from classical mechanics, and any initial state will rapidly decohere and lose all apparent memory of initial conditions. Local observables, for example, will acquire expectation values in agreement with conventional statistical mechanical ensembles. Such systems are labelled “chaotic” or “quantum ergodic”, and these are well explained by the celebrated Eigenstate Thermalisation Hypothesis (ETH) [3–6], which asserts that the eigenstates of the Hamiltonian themselves resemble the appropriate thermodynamic ensemble, under the lens of any local measurement. However, there are certain ways to avoid this fate, violate the ETH, and remain far from equilibrium. Integrable systems have an extensive number of exactly conserved quantities which constrain the dynamics, but sit at finely-tuned points in parameter space and are therefore unstable to any arbitrary perturbation. In many-body localisation (MBL), on the other hand,

the application of strong quenched disorder leads to exponential localisation of the energy eigenstates in the Hilbert space and the arrest of transport in a wide variety of systems, robustly preventing equilibration [7, 8].

In both integrable and MBL systems, every single eigenstate violates the ETH, a situation known as strong ergodicity breaking. But the experimental discovery of many-body quantum scars, eigenstates of an otherwise chaotic Hamiltonian which are non-thermal and can lead to anomalous dynamics including periodic “revivals”, in Rydberg atom simulators [9] has shown that there is still much missing from our understanding of quantum thermalisation. Such “scarred” systems sit in-between fully localised and fully chaotic, with the dynamics dependent on initial conditions and only a vanishing fraction of states violating the ETH in the thermodynamic limit: this is known as weak ergodicity breaking. While recent theoretical developments have connected certain such scars to “towers of states” resulting from broken symmetries [10–16], there are also scars left unexplained by this paradigm [16–19], and we are far from a complete and unified picture [20].

In contrast to the problem of thermalisation, which is generally concerned with the behaviour of systems at finite or even *infinite* temperature, there is the seemingly unconnected phenomenon of topological phases of matter, in which systems without local order parameters or spontaneous symmetry breaking exhibit distinct phases of matter at *zero* temperature. In particular, we will be concerned with symmetry-protected topological (SPT) order [21–27], which is usually associated with protected edge modes, degenerate entanglement spectra, and string order. Both topological order and SPT order may also be exploited to encode quantum information, such as in the edge modes [28]. These phenomena are protected at zero temperature as long as the ground state remains gapped and perturbations to the Hamiltonian respect a certain symmetry, but at finite temperature, interactions with the bulk can cause the edge modes to decohere, and other signatures of topological order are lost. Hence, topological order is generally only a property of the ground state.

Despite this limitation, there have been various approaches to preserving topological order at finite energy density. Since MBL and quantum scarring can both arrest the approach to thermal equilibrium, contrary to the ergodic dynamics expected of generic quantum systems, this raises the possibility that by combining SPT order

with such ergodicity breaking phenomena one might be able to construct a quantum memory that is robust at finite temperature. In particular it has been shown that MBL, induced by strong disorder, can protect topological edge modes and other signatures of topological order, up to arbitrarily high energies [29, 30]. The question therefore remains whether other non-ergodic phenomena, such as quantum scarring, would be able to preserve edge modes at finite temperature without destroying transport. Furthermore, a full picture of the interactions between topological order and MBL is still developing. Whereas previously it was thought that localisation could be preserved through topological phase transitions [31], recent evidence from small-system studies has suggested this may not be possible [32–35].

The existence of an MBL phase has been rigorously proved in one dimension given certain physically reasonable assumptions [36]. Yet in recent years, MBL has come under renewed scrutiny, with the discovery that rare low-disorder regions which thermalise locally can then destabilise the entire localised phase in a “thermal avalanche” driven by many-body resonances [37–45]. In fact, it is even argued by some that the MBL phase does not actually exist in physically realistic systems, instead only surviving as a finite-lifetime or finite-size regime. Understanding the precise nature of these avalanches, the role resonances play in the delocalisation transition, and the strength of disorder necessary to ensure a localised phase in the thermodynamic limit are therefore all key challenges facing the MBL community.

The remainder of this thesis is arranged as follows. In Chapter 2, we cover the necessary background in detail, looking at thermalisation, MBL, quantum scarring, and topological order. Then, in Chapter 3, we study a certain spin-1/2 chain which hosts $\mathbb{Z}_2 \times \mathbb{Z}_2$ SPT order, resulting in a long-lived prethermal edge mode when the even and odd sublattices are detuned. We show that quantum many-body scars may be responsible for this behaviour, and furthermore that these scars host signatures of SPT order. Chapter 4 explores the infinite-temperature phase transition between MBL phases with different SPT orders, which has gained attention recently due to small-system studies which show an ergodic phase intervenes at this transition [32–35]. Using a real-space renormalisation group technique, we characterise the localised phases and uncover many-body resonances between the localised states, showing that these proliferate near the critical point and destroy the localised phase.

Using this same technique, as well as applying perturbation theory to free-fermion systems, in Chapter 5 we look at a disordered spin-1/2 Ising chain with transverse and longitudinal fields (as studied in Ref. [36]’s proof of MBL), finding evidence that proximity to an SPT critical point destabilises the MBL phase, even though the model itself has no symmetries. Finally, in Chapter 6, I summarise this thesis and present my conclusions.

Chapter 2

Background

2.1 Thermalisation and Quantum Chaos

In classical statistical mechanics, no matter how complex the initial state, a large interacting system will eventually reach an equilibrium state described only by macroscopic conserved quantities such as energy and particle number. It is possible to slow this process down, for example with kinetic barriers or by lowering the temperature, but in the infinite-time limit this fate is impossible to avoid. Likewise, when initialised in a highly-localised pure state, a generic closed quantum system will rapidly decohere and lose all apparent memory of initial conditions [5]. Furthermore, it will explore every part of the Hilbert space, such that the long-time average is described by standard statistical mechanics; such a system is said to be *ergodic*. Indeed, the classical case is simply the high-energy limit of the quantum case, as specified by the correspondence principle, and so both must exhibit similar behaviour. This presents a paradox, however. Quantum mechanics is governed by unitary time evolution, which implies that in a closed system no information is lost and that a pure state will always remain a pure state. The solution comes from considering a classical analogy: if, after reaching equilibrium, we were able to record the position and velocity of every particle in our system, by reversing the direction of motion we would be able to play the physical evolution backwards in time and so reconstruct the initial conditions. However, to do so would clearly require infeasibly many measurements. The initial conditions were never lost – they just became inaccessible.

The same idea applies to our quantum system. Under unitary time evolution, subsystems act as heat baths to each other, scrambling the degrees of freedom and spreading information throughout the system. In particular, the physical degrees of

freedom become highly entangled with each other, leading to a “volume law” regime in which bipartite entanglement entropy is proportional to the volume of the system. As a result, only extensive global measurements can recover the initial conditions – even though it remains in a pure state for the entire evolution, for all intents and purposes, the system behaves as if it has reached thermal equilibrium.

This behaviour should be compared to that of an integrable system, in which the dynamics are constrained by an extensive number of conserved quantities known as integrals of motion [46, 47] – or equivalently, which can be mapped to a system of non-interacting particles. In this case, memory of initial conditions is preserved for all time via the integrals of motion, and there is no meaningful approach to thermal equilibrium. However, these systems are unstable to generic perturbations, and so do not represent a true phase of matter, but only a finely-tuned special case.

2.1.1 The Eigenstate Thermalisation Hypothesis

The above intuition is formalised in the Eigenstate Thermalisation Hypothesis (ETH) [3–6, 48, 49], which in simple terms asserts that eigenstates of the system resemble thermal ensembles, and is inspired by random matrix theory which was originally developed to understand the excited eigenstates of heavy nuclei [5]. More precisely, it is an ansatz for the matrix elements of local, physically reasonable operators¹ which is expected to be obeyed by a generic non-integrable quantum system, and which numerical evidence has shown to be the case for a wide variety of such systems [51–59]. It states that,

$$\langle m | \mathcal{O} | n \rangle = \mathcal{O}(\bar{E}) \delta_{mn} + e^{-S(\bar{E})/2} f_{\mathcal{O}}(\bar{E}, \omega) R_{mn} , \quad (2.1)$$

where $|n\rangle$, $|m\rangle$ are energy eigenstates of the system, with $\bar{E}$ and ω their average energy and energy separation respectively, $S(E)$ is the thermodynamic entropy at energy E , and R_{mn} is randomly distributed with zero mean and unit variance. (R_{mn} and $f_{\mathcal{O}}(\bar{E}, \omega)$ are real when time-reversal symmetry is present, otherwise complex; additionally, R is Hermitian.) Additionally, and crucially, both $\mathcal{O}(\bar{E})$ and $f_{\mathcal{O}}(\bar{E}, \omega)$ are smooth functions of their arguments. The smoothness of $\mathcal{O}(\bar{E})$ in the first

¹Generally, it is expected to be obeyed by generic few-body operators, even those with support up to half the system size [50]. On the other hand, operators such as eigenstate projectors obviously violate this ansatz.

term ensures that the expectation value in a pure state matches the microcanonical ensemble average, because in any narrow energy window the diagonal terms will be near-constant. Because thermodynamic entropy is extensive, fluctuations away from this due to the second term are then exponentially suppressed in system size, such that local measurements of a pure state in a system obeying the ETH look “thermal”.

To elaborate on this statement, consider now the density matrix of a pure state $|\psi(t)\rangle = \sum_n c_n |n\rangle$ as it evolves in time, in the energy eigenbasis $\{|n\rangle\}$. On the diagonal, we have the eigenstate decomposition $\rho_{nn}(t) = |c_n|^2$, while the off-diagonals describe the coherence between the different components, $\rho_{nm}(t) = c_n c_m^* e^{i(E_n - E_m)t}$. It is these off-diagonals therefore that encode the precise initial conditions. However, since the energies of states are in general irrational multiples of each other, time evolution gives these off-diagonal terms random phases – scrambling the coherences such that the long-time average $\frac{1}{T} \int_0^T \rho(t) dt$ tends to a matrix containing only the diagonal terms, the so-called diagonal ensemble ρ_{DE} . As a result, observers at long times record a classical mixture of the component eigenstates, without quantum coherence, and likewise local measurements return simple weighted averages of the eigenstate expectation values. By guaranteeing that these expectation values match the microcanonical ensemble, the ETH therefore ensures macroscopically large systems behave exactly as thermodynamics predicts and eventually thermalise [5]. One may note that the mean level spacing is exponentially small in system size, and therefore some off-diagonal terms may have exponentially *long* lifetimes, leading to memory of initial conditions at intermediate times. However, the weighting $e^{-S(\bar{E})/2}$ in the second term of the ETH ensures the contribution of these terms to expectation values remains small, enabling rapid approach to equilibrium².

When focusing on a narrow energy window such that the spectral function $f_{\mathcal{O}}(\bar{E}, \omega)$ can be treated as constant, the second term ensures the ansatz matches the predictions of random matrix theory. For diffusive systems, this generally requires that ω is less than the Thouless energy $E_T = D/L^2 = 1/\tau_D$, where τ_D is the diffusion time, which, despite vanishing in the thermodynamic limit, will still contain an exponentially large number of energy levels. Otherwise, this threshold is given by

²In fact, $e^{-S(\bar{E})}$ is exactly the level spacing (in appropriate units).

$1/\tau^*$, with τ^* the longest timescale in the system [5]. (There is some disagreement; Refs. [59] find a much tighter bound on the energy scale.) On the other hand, at large energy separations, $f_{\mathcal{O}}(\bar{E}, \omega)$ decays exponentially in ω . The spectral function can be directly connected to two-time thermal correlation functions through the Kubo-Martin-Schwinger (KMS) relation [60, 61] as,

$$\langle \mathcal{O}(t) \mathcal{O}(0) \rangle_{\beta} - \langle \mathcal{O} \rangle_{\beta}^2 = k_2^{\beta}(t) = \int |f_{\mathcal{O}}(E(\beta), \omega)|^2 e^{-\beta\omega/2} e^{i\omega t} d\omega, \quad (2.2)$$

where $\langle \bullet \rangle_{\beta}$ is the thermal expectation value at inverse temperature β and $E(\beta)$ the thermodynamic energy at said temperature, and therefore provides important insights into the dynamics of a system. $f_{\mathcal{O}}(\bar{E}, \omega)$ also gives the response of a system to an external perturbation that pushes it away from equilibrium [62].

While the ansatz (2.1) assumes independence between the matrix $\langle m | \mathcal{O} | n \rangle$ elements, this is not usually true in practice. Recent work has therefore extended the ETH ansatz to capture correlations between these with a set of generalised spectral functions, which then lead directly to an expansion of multi-time thermal correlation functions in terms of so-called thermal free cumulants [63, 64].

In fact, there are both strong and weak versions of the ETH [6]. The strong ETH asserts that *every* eigenstate at finite energy density (within each symmetry sector) obeys the ETH ansatz, and a generic interacting many-body system usually satisfies this hypothesis. However, the weak ETH allows for finitely many states, or a vanishing proportion in the thermodynamic limit, to violate the ansatz. While integrable and many-body localised systems violate both forms of the ETH, many-body quantum scars only violate the weak ETH. These cases are discussed in subsequent sections.

2.1.2 Level Statistics

The ETH ansatz is inspired by the observation that, when limited to small energy windows, the Hamiltonians of chaotic complex systems resemble random matrices whose entries only depend on the symmetries at hand. Specifically, following the ideas of random matrix theory, after resolving symmetries one can replace these Hamiltonians by random matrices with normally-distributed entries, drawn from ensembles which are invariant under either orthogonal or unitary transformations. These ensembles are known as the Gaussian Orthogonal Ensemble (GOE) and

Gaussian Unitary Ensemble (GUE), respectively, and correspond to systems with and without time reversal invariance [65, 66]. Equivalently, GOE and GUE matrices have real and complex entries respectively. (There is in fact one further ensemble with quaternionic entries, the Gaussian Symplectic Ensemble or GSE, but it is less common). Chaotic many-body Hamiltonians therefore exhibit *universal* properties, one of which is their energy level statistics. This may be measured by the $\tilde{r}$ -statistic [67, 68],

$$\tilde{r}_j = \min \left(\frac{E_{j+1} - E_j}{E_j - E_{j-1}}, \frac{E_j - E_{j-1}}{E_{j+1} - E_j} \right), \quad (2.3)$$

where E_j is the j^{th} energy level. For chaotic Hamiltonians, hybridisation between quasiparticle modes due to interactions leads to level repulsion, which ensures that $P(\tilde{r}_j = 0) = 0$. The full form of the probability distribution then depends on whether the Hamiltonian follows GOE or GUE statistics, but is universal for systems of a particular class.

On the other hand, in large integrable systems the energy levels are uncorrelated and follow a Poisson distribution – crucially, with $P(\tilde{r}_j = 0) > 0$ and no level repulsion. Hence the level statistics of many-body Hamiltonians are a robust numerical test for quantum chaos.

2.1.3 The Quantum Information Perspective

As alluded to above, ideas from quantum information provide deep insight into the nature of thermalisation. Instead of analysing the flow of matter and energy throughout a system, in this picture we look at the spread of information – specifically, how entanglement takes the information encoded by the initial conditions, and distributes it over the physical degrees of freedom such that it becomes unrecoverable by any experimentally-reasonable process. In this way, memory of the past is lost, and time’s arrow applies equally well to quantum systems as it does to classical ones.

A key quantity of interest to the field of quantum information is *entanglement entropy*. Suppose we divide a system into two subsystems A and B with $\dim A \leq \dim B$, such that the total Hilbert space is given by the tensor product of the Hilbert spaces of the respective subsystems, and we may cleanly partition degrees of freedom between the two. Suppose furthermore that this system is in some pure state $|\psi\rangle$, such that the density matrix for the system is given by $\rho_{AB} = |\psi\rangle\langle\psi|$. We then define the *reduced* density matrix $\rho_A = \text{Tr}_B\{\rho_{AB}\}$, where we trace over the degrees




Figure 2.1: Under unitary time evolution from an initial pure state (top left), the degrees of freedom of a generic quantum system are strongly mixed, resulting in all-to-all entanglement and the spreading of information throughout the system (bottom left). Even though the system remains in a pure state throughout, any local measurement cannot recover information about the initial state, and as such the expectation values and distribution of measurement outcomes looks “thermal” (bottom right). This is to be compared with local measurements immediately after the quench, which give outcomes far from the equilibrium distribution (top right). Likewise, the reduced density matrix of any subsystem cuts through entanglement proportional to the volume of that subsystem, and we have a volume-law entanglement entropy. [From Ref. [70]. Reprinted with permission from AAAS.]

of freedom in B . If the degrees of freedom in A are entangled with those in B , then after tracing out the latter we cannot know the state of A exactly, and so while ρ_{AB} is a pure state, ρ_A is a mixed state. The bipartite entanglement entropy is then defined as,

$$S = -\text{Tr}\{\rho_A \ln \rho_A\} . \quad (2.4)$$

If the total state has some Schmidt decomposition $|\psi\rangle = \sum_j \sqrt{\lambda_j} |j\rangle_A |j\rangle_B$, then $\rho_A = \sum_j \lambda_j |j\rangle_A \langle j|_A$, and so $S = -\sum_j \lambda_j \ln \lambda_j$. Therefore, when $|\psi\rangle$ is separable, $S = 0$, and otherwise grows as the two subsystems become more entangled. Intuitively, S measures how much entanglement we need to “cut” to separate the two subsystems: if there are n coefficients $\lambda_j = 1/n$, then $S = \ln n$ (and less if the weight is concentrated among few coefficients). Due to this connection, the Schmidt coefficients λ_j are also referred to as the entanglement spectrum. Remarkably, despite this being a purely quantum mechanical quantity, it turns out that the entanglement entropy coincides with the thermodynamic entropy at equilibrium [69]. This suggests a deep connection between entanglement and thermalisation.

This connection comes from considering what happens to the degrees of freedom as they evolve under a chaotic many-body Hamiltonian. Interactions mix the degrees of freedom, generating entanglement between them, and given enough time this spreads throughout the system: each degree of freedom is entangled with every other degree of freedom. Now, any local measurement only has access to a partial set of the degrees of freedom, and cannot recover the full (pure) state, instead seeing a classically mixed state. Because of this all-to-all entanglement pattern, the entanglement entropy will scale extensively with the volume of the subsystem, $S \propto V$ (see Fig. 2.1). States with this scaling are said to be *volume law* [5, 70, 71].

On the other hand, consider the ground state of a gapped Hamiltonian. Unlike a state at high energy, which is in the bulk of the spectrum and so easily couples to many nearby states, such states are isolated and are therefore special, with atypical expectation values and short-range entanglement³. When we cut our system into two subsystems A and B , the bipartite entanglement entropy will only be proportional to the size of the boundary, $S \propto \partial V$ – a so-called *area law*. Ground states near criticality, where the gap closes, may also exhibit a logarithmic correction, $S \propto \log V$ [72]. Integrable systems, with extensively many integrals of motion, also typically exhibit a broad distribution of entanglement entropies, including both area and volume law states. Entanglement is therefore a key diagnostic of chaotic vs integrable dynamics.

2.2 Avoiding Thermalisation: Many Body Localisation

We have covered thermalisation, the Eigenstate Thermalisation Hypothesis, and how systems which satisfy the latter will rapidly approach a thermal equilibrium state with no memory of initial conditions. However, there are certain known classes of quantum system which violate the ETH and avoid thermalisation entirely. One of these is the class of finely-tuned integrable models – systems with an exactly solvable spectrum, in which an extensive number of conserved quantities known as integrals of motion constrain the dynamics [5]. For example, the transverse field Ising model can be mapped to free fermions through a Jordan-Wigner transform. In such a system, every single eigenstate violates the ETH, and memory of initial conditions is

³Topological order (see Sec. 2.4) can lead to long-range entanglement in the ground state, but this leads to a constant term in the entanglement entropy.

preserved for all time, but a generic perturbation is enough to destroy integrability, leading to chaotic dynamics and a rapid approach to thermal equilibrium. Naturally, this led to the question of whether there exists a way to more robustly avoid the fate of thermalisation. One such promising mechanism comes from the fascinating phenomenon of localisation.

2.2.1 Anderson localisation

Consider a particle on a crystal lattice, with a short-ranged hopping term and a site-dependent on-site energy. For concreteness, we can write the free-fermion Hamiltonian,

$$H = \sum_{\mathbf{j} \in \Lambda} \varepsilon_{\mathbf{j}} c_{\mathbf{j}}^{\dagger} c_{\mathbf{j}} + \sum_{\mathbf{j}, \mathbf{k} \in \Lambda} V_{\mathbf{j}\mathbf{k}} c_{\mathbf{j}}^{\dagger} c_{\mathbf{k}} + \text{h.c.} , \quad (2.5)$$

for some lattice Λ . Here, $c_{\mathbf{j}}^{\dagger}$ ($c_{\mathbf{j}}$) is the creation (annihilation) operator at vector-indexed site $\mathbf{j}$. In the case of translational invariance, Bloch's theorem tells us that the single-particle eigenstates will be extended plane waves, and from such a model we would expect to see diffusive or ballistic transport throughout the entire lattice [73, 74]. However, if we instead allow the on-site energies $\{\varepsilon_{\mathbf{j}}\}$ to be position-dependent (but static), picking them randomly from some probability distribution with width W , remarkably the eigenstates become exponentially localised in space – even for arbitrarily small W in one and two dimensions. (In three dimensions, W must be larger than the hopping energy scale). This remains true even if the hopping is long-range: the only requirement is that $V_{\mathbf{j}\mathbf{k}}$ falls off faster than $1/|\mathbf{j} - \mathbf{k}|^3$. This has dramatic consequences for the dynamics: a wavefunction initialised at a particular location in space will remain within an exponential envelope with some localisation length ξ , with a finite probability density at the starting location. To be precise, if the wavefunction at $t = 0$ is given by $\psi_{\mathbf{j}}(t = 0) = \delta_{\mathbf{j}0}$, then the long-time asymptotic form at long distances is given by $\psi_{\mathbf{j}}(t \rightarrow \infty) \sim \exp(-|\mathbf{j}|/\xi)$, and the energy eigenfunctions have a similar form. This phenomenon is now known as Anderson localisation, after P.W. Anderson [75].

One could compare Anderson localisation to a ball rolling around a hilly landscape, whose motion is impeded by the peaks and troughs in the way. But this naive classical picture is not accurate – for sufficient disorder, the entire spectrum becomes localised, even when the energy of highly-excited states exceeds the disorder width.

Additionally, a quantum ball would eventually tunnel through energy barriers in its way, whereas Anderson localisation survives even in the infinite-time limit. The correct approach, used by Anderson in his 1958 paper [75], is to start from the infinite disorder case (in which particles are trivially localised) and then perturbatively expand the Green's function around this point in powers of the hopping term. This is (today) known as the locator expansion, and is equivalent to summing over the various paths through the lattice that a particle can take to reach a particular destination. The contributions from these paths destructively interfere, preventing the particle from fully exploring the lattice. One can also use this approach to calculate the amplitude at the original site – which is only non-zero in the infinite-lattice limit if the particle is indeed localised. The requirement to account for loops makes these calculations challenging, but simplified models on e.g. Cayley trees, which do not contain loops, allow for exact results [76].

The existence of Anderson localisation for arbitrarily weak disorder has been proved rigorously in one dimension [77], while strong arguments – based on a scaling theory of localisation and conductance – have been made for the same in two dimensions [78]. The phenomenon has also been verified experimentally in a diverse range of media, including cases with practical applications such in optical fibres [79–85].

There are some interesting connections between Anderson localisation and certain other phenomena, both classical and quantum. A random walk in one or two dimensions always returns to the origin, whereas a walk in 3D may not; likewise, an arbitrarily small attractive potential will lead to bound states in 1D or 2D [86], but this is not the case in 3D. This emphasises the importance of dimensionality in the problem of localisation. Furthermore, for a weakly disordered potential in 1D, the localisation length is equal to the mean free path derived due to backward scattering of waves in a classical kinetic theory [87]. In fact, Anderson localisation isn't even a strictly quantum phenomenon, but a property of ordinary waves: for example, sound waves have been localised in a disordered elastic network [88].

Finally, note that Anderson localisation should not be confused with Mott localisation, which occurs due to strong repulsive interactions rather than disorder and does not preclude thermal equilibrium; or weak localisation, a precursor to

Anderson localisation occurring at weaker disorder strength, in which paths that loop back on themselves constructively interfere leading to a (finite) increase in resistivity.

2.2.2 The interacting many-body case

Even though arbitrarily weak disorder is enough to localise a single particle in fewer than three dimensions, we do not observe localisation in our every day lives. Why? For starters, real-world systems are not closed, like the model (2.5), but instead open systems which can exchange energy and particles with their environment. It has been shown that coupling an Anderson localised system to a heat bath delocalises it – as a rough picture, by exchanging energy with the heat bath, the particle can move between localised states that are not exactly degenerate, and so travel through the system in a process known as variable range hopping [89]. Moreover, even when closed, most real physical systems are interacting and their quasiparticles can exchange energy with each other, which likewise has the potential to delocalise particles and destroy memory of initial conditions. Another perspective would be to consider the many-body case as a single particle hopping around a lattice representing the Fock space and its connectivity under the Hamiltonian. Because there are far fewer parameters in a many-body Hamiltonian than states in the exponentially large Fock space, any disorder in the matrix elements (which give the effective hopping strengths) would necessarily be strongly correlated, precluding the usual arguments for Anderson localisation.

For many years therefore it was an open question whether localisation could survive in the presence of interactions, a phenomenon now termed many-body localisation (MBL). The existence of such a phase, robustly violating the ETH and able to avoid thermalisation, would have substantial ramifications for conventional statistical mechanics; we cover these consequences in Sec. 2.2.3. This was, and still is, a challenging problem: interacting systems are vastly more complicated to simulate than free systems, and exhibit qualitatively different phenomena. It is only in the last two decades, especially with the rise of powerful general-purpose computing, that we have started to gain real insight into MBL. Despite this, there are still many open questions, and some controversy too, including some who now argue that the critical disorder strength is infinite in the thermodynamic limit – see Sec. 2.2.4 for details.




Figure 2.2: Eigenstate-resolved spin-spin connected correlation $\langle \sigma_i^z \sigma_{i+d}^z \rangle_c$ against separation distance d , averaged over disorder and eigenstates, for various system sizes $L = 8, 10, \dots, 16$ and disorder widths h in a random-field Heisenberg model. For $h = 6.0$, a clear exponential decay is shown for the correlation, indicating localisation. [Reprinted figure with permission from Ref. [90]. Copyright (2010) by the American Physical Society.]

Theory suggested that inelastic collisions would lead to delocalisation through dephasing, but raised the possibility that elastic processes would not necessarily have the same consequences, thus permitting a localised phase at finite temperature [91–93]. This prompted a slew of theoretical [94, 95] and numerical [67, 90, 96] investigations into MBL as a distinct phenomenon from Anderson localisation, inspired by these developments. Ref. [90] looked at the level statistics (2.3), exponential drop-off of spin-spin correlation in the energy eigenstates (Fig. 2.2), memory of initial conditions, and violation for the smoothness of $\mathcal{O}(\bar{E})$ in the ETH (2.1) to prove the existence of an MBL regime in the spin-1/2 Heisenberg chain with random fields; meanwhile Ref. [96] used entanglement growth (2.4) in tensor network simulations of a disordered XXZ chain, showing that this saturated over time.

The consensus that emerged was that localisation could survive in the many-body setting if the disorder was stronger than the interactions [7, 8]. For example, Ref. [90] estimated the critical disorder width to be about 3-4 times the interaction strength (although different measures were not in agreement, and also showed finite-size drift). Indeed, as predicted by some of those papers, programmable quantum simulators soon




Figure 2.3: (a) The equilibrium state for cold rubidium atoms in a 2D optical lattice is circular, due to a rotationally symmetric confining potential. Quenched disorder may be applied to this lattice by means of a digital mirror device. (b) When no disorder is applied and the atoms are taken out of equilibrium (by removing half to form a semi-circle), they rapidly thermalise and return to equilibrium. (c) However, in the presence of sufficiently strong quenched disorder, the atoms retain some memory of their initial condition for all time; in fact, the imbalance between atoms on the left (L) and right (R), $\mathcal{I} = \frac{N_L - N_R}{N_L + N_R}$, tends to a non-zero value (see reference for details). [From Ref. [97]. Reprinted with permission from AAAS.]

showed strong experimental evidence for the phenomenon [97–99]. Fig. 2.3, taken from Ref. [97], provides a particularly striking example of MBL in a two-dimensional cold atom system, specifically Rubidium atoms in an optical lattice. The atoms were arranged in a highly out-of-equilibrium state, confined to the left-hand side of the system, and then allowed to hop through the lattice under the influence of varying levels of quenched disorder. Remarkably, when the disorder was strong enough, the atoms remained almost entirely localised to the same side of the lattice: they retained memory of initial conditions. Analysis of the imbalance between atoms on the left (L) and right (R) of the system, $\mathcal{I} = \frac{N_L - N_R}{N_L + N_R}$, showed this quantity tending to a non-zero value, within the experimentally accessible timescale, providing a convincing proof of many-body localisation. The experiment was also able to show that the critical disorder depended on the interaction strength, by moderating the density of atoms. Another experiment, shown in Ref. [98], placed atoms only on the even sites of a one-dimensional system; with sufficient disorder, these too retained memory of this initial condition.

However, these approaches suffer drawbacks. The theoretical approaches relied on strong assumptions or approximations. Numerics are necessarily limited to finite system sizes, particularly in the case of exact diagonalisation where the state of the art is $L \simeq 20$, and the presence of disorder prevents easy extrapolation to the thermodynamic limit. Experiments, on the other hand, are often limited to short timescales due to the accumulation of errors and interactions with the environment. Even with these problems resolved, Ref. [100] claims that the microscopic timescales necessary for a system to exhibit true many-body behaviour will always exceed the Heisenberg time (after which a finite-size system exhibits quasiperiodic dynamics) for achievable system sizes in experiment. A breakthrough came with J.Z. Imbrie’s celebrated proof of MBL in 1D [36], in which the author shows that “many-body localisation follows from a physically reasonable assumption that limits the amount of level attraction in the system”. The work focuses on an Ising model with longitudinal and transverse fields,

$$H_{\text{Imbrie}} = \sum_j h_j \sigma_j^z + \gamma_j \sigma_j^x + J_j \sigma_j^z \sigma_{j+1}^z , \quad (2.6)$$

where $\gamma_j = \gamma \Gamma_j$ with $\gamma \ll 1$, and h_j, Γ_j, J_j are smaller than 1 and drawn from independent random distributions with bounded probability density. The author was able to show that the model could be diagonalised through a series of quasi-local unitary transformations so long as the model did not exhibit strong level attraction. Given that chaotic models show level repulsion and integrable models show neutral (that is, Poisson) level statistics, this seemed a reasonable assumption, and allowed Imbrie to deal with strong resonances that could otherwise not be dealt with perturbatively. Since the transformations were quasi-local – meaning that spin flips over large distances were exponentially suppressed – the model must be in the MBL phase.

The proof only holds validity in the regime of very small γ , but regardless is strong evidence of MBL as a robust phase of matter in the thermodynamic limit. Chapter 5 explores the full phase diagram of the model, allowing for large γ . However, as will be explored in Sec. 2.2.4, there are many recent signs that MBL may not be as stable as previously thought. But before we get onto that, in the next subsection we will cover the various key signatures and consequences of MBL, what these mean

for statistical mechanics, and the possible real-world applications of MBL including to quantum technologies.

2.2.3 Consequences of many-body localisation

By violating the key assumptions of conventional statistical mechanics – namely, that systems will fully explore phase space up to the limits imposed by symmetries and the resultant conservation laws – MBL has stark consequences for the behaviour of quantum systems. For starters, it is clear that if memory of initial conditions is retained for all time, then systems cannot relax to an equilibrium state determined only by those aforementioned macroscopic conserved quantities. This was already true for Anderson localisation – but this only held in the limit of zero interactions, while robustness to integrability-breaking perturbations would purport to make MBL a true phase of matter. Despite being a phase of matter, the transition from ETH to MBL can only be detected by measuring dynamical quantities – any average over a thermal ensemble of states will produce results corresponding to thermal equilibrium [7]. Additionally, the unique out-of-equilibrium nature of MBL allows it to bypass a number of results from conventional statistical mechanics. For example, the Landau-Peierls theorem prohibits ordering and spontaneous symmetry breaking at finite temperature in 1D quantum systems at thermal equilibrium, due to thermal fluctuations which intrinsically destroy said ordering. But of course MBL is not at thermal equilibrium, and thus is able to exhibit this forbidden order even in 1D [7, 31].

Many-body localisation was initially understood as an extension of Anderson localisation, or as Anderson localisation in a many-body phase space, and interpreted in that context. This means that early studies focused on, e.g., frozen dynamics and the lack of DC transport. But today, MBL is usually thought of as a form of *emergent* integrability, controlled by a set of quasi-local integrals of motion termed “l-bits” which, too, are emergent [101, 102]. Focusing on the spin-1/2 case, take a simple non-interacting model in which the energy eigenstates are product states in the σ^z basis, such as (2.6) with $\gamma = 0$. $[H, \sigma_j^z] = 0$ for all j , and so these spins are integrals of motion. If we then add weak interactions, such that the system remains MBL, then the new eigenstates will still be close to product states, and in fact we will be able to uniquely identify them by a new extensive set of integrals

of motion $\tau_j^z = U^\dagger \sigma_j^z U$, such that $[H, \tau_j^z] = 0$. Here, the unitary transformation U only creates spin flips between nearby degrees of freedom, with long-range flips exponentially suppressed, such that the τ_j^z themselves have exponentially decaying tails – a property known as quasi-locality. We can then write down a diagonalised Hamiltonian as,

$$U^\dagger H U = H_{\text{diag}} = \sum_i \varepsilon_i^{(1)} \tau_i^z + \sum_{ij} \varepsilon_{ij}^{(2)} \tau_i^z \tau_j^z + \sum_{ijk} \varepsilon_{ijk}^{(3)} \tau_i^z \tau_j^z \tau_k^z + \dots, \quad (2.7)$$

where in general couplings between distant l-bits are suppressed. Furthermore, one can also define a corresponding set $\tau_j^x = U^\dagger \sigma_j^x U$, thus providing a complete generator set in operator space. Because they are constructed by a unitary operator from the bare spins, the l-bits obey all the usual spin commutation relationships, and a many-body eigenstate is specified uniquely by the eigenvalues of the l-bits. While constructing the transformation U is non-trivial in most cases, Refs. [36] prove the existence of this transformation, and therefore the existence of l-bits, in one dimension (given the assumption of limited level attraction). However certain numerical techniques, such as those based on the strong-disorder renormalisation group [31, 103–106], flow equations [107–109], tensor networks [110–114], or otherwise [115–117] have had substantial success in approximating l-bits for a variety of systems, particularly in the perturbative limit of strong disorder and weak interactions – where the l-bits may be compared to the exactly solvable integrals of motion found at a nearby (disordered) integrable point.

The phenomenology of MBL is varied and complex, but there are a number of key features generally agreed upon, and used to detect MBL in numerics and experiment, which include:

- Localisation in Hilbert space
- Absence of DC transport
- Area-law entanglement of energy eigenstates
- Slow (logarithmic) growth of entanglement following a quench
- Slow (power-law) relaxation of local observables to non-equilibrium values, and hence a failure to approach equilibrium in general.
- Strong violation of the ETH

- Poisson level statistics (i.e. no level repulsion)

These all neatly follow from the existence of l-bits [102]. Given the exponential decay of l-bit support, any eigenstate of the system (which is itself a mutual eigenstate of all of the l-bits) will be localised in Hilbert space. Unlike in clean integrable systems, which can exhibit ballistic or diffusive transport, the quasi-local nature of the l-bits ensures that long-range transport (beyond the localisation length) is suppressed and that MBL systems are insulators. A rigorous argument is given in Ref. [118]. Area-law entanglement, where the entanglement entropy is proportional to the area of a cut (so constant in 1D, where cuts are 0D points) follows for a similar reason – only l-bits with centres within the localisation distance of the cut contribute substantially to the entanglement. Since the energy eigenstates are simultaneous eigenstates of the l-bits, violation of the ETH in every eigenstate follows immediately as the l-bits should obey the ansatz. Strictly local operators are also quasilocal deformations of the l-bits (by definition), and these too will violate the ETH; likewise, because the l-bits are conserved in the dynamics, the expectation value of local operators will not relax to equilibrium. The lack of level repulsion follows from the fact that nearby energy eigenstates will have very different l-bit configurations, preventing them from hybridising.

Finally, we should explain the slow dynamics following a quench from out of equilibrium. Because the interactions between l-bits [the higher order terms in Eq. (2.7)] fall off exponentially with l-bit separation, dephasing is exponentially slow between well-separated pairs of l-bits. This means, conversely, that the distance over which dephasing acts grows *logarithmically* in time, and likewise the number of l-bits contributing to entanglement [119, 120]. However, unlike with the energy eigenstates, the entanglement following a quench will grow to a limit that scales with the volume of the system (and is unbounded in the thermodynamic limit). A related argument is made in Ref. [121] for the power-law relaxation of local observables. This dephasing makes MBL qualitatively different from Anderson localisation, in which there are no interactions between localised particles and so entanglement growth following a quench is bounded.

The area-law entanglement entropy of eigenstates is particularly striking, as this is a property that MBL states share with the ground states of gapped Hamiltonians.

The l-bit picture makes this clearer: any product state in the l-bit basis is also the ground state of a quasi-local, gapped Hamiltonian, so in some sense MBL eigenstates resemble many-body ground states despite being highly excited. The area-law entanglement also means it is possible to accurately approximate MBL eigenstates using tensor networks, such as matrix product states [MPS, see Eq. (2.11)] and projected entangled-pair states (PEPS), thus making MBL amenable to efficient simulations on classical computers such as density matrix renormalisation group (DMRG) or time-evolving block decimation (TEBD) [122]. The entanglement spectrum (2.4) has also been shown to have a distinctive power-law form in MBL systems, due to the l-bit structure of the eigenstates [123].

MBL has been proposed as a form of quantum memory [124], for example in quantum computers which currently face challenges in keeping qubits coherent over macroscopic timescales. To reliably store quantum information, both the z - and x -components of the l-bit must be preserved (equivalent to keeping the l-bits in phase with each other). However, while the z -components are (by definition) conserved by time evolution, the l-bits will precess about the z -axis due to interactions with the other l-bits, leading to the dephasing mentioned above and a changing x -component [7, 125]. Hence, only classical information is preserved. In principle, it is still possible to use a modified spin-echo protocol to store quantum information in an l-bit, and then read it out at some later time – but crucially, this only works for a single qubit. To store multiple qubits would either take an exponentially complicated protocol, thus entirely defeating the point of a quantum memory [126, 127], or inefficiently require large spatial separation between the qubits. This means that MBL would have to be combined with other phenomena, such as topological order, to enable efficient preservation of quantum information [29, 30].

2.2.4 Breakdown of many-body localisation

There are some subtleties in the picture of MBL painted in the previous section. Ref. [115] claims that, in two or more dimensions, the l-bits become susceptible to boundary instabilities and become “l*-bits”, operators which are only *approximately* conserved with $[H, \tau_j^{*z}] \sim O(\exp(-L/\xi))$. This is reminiscent of strong zero modes [128], which display similar scaling. In the thermodynamic limit, the dynamics is frozen, but at finite sizes these decay exponentially slowly. Furthermore, the l*-bits

obey the ETH, but with very narrow spectral functions. With this phenomenology, many of the defining features of MBL would vanish, including area-law entanglement, lack of level repulsion, and the ability to sustain forbidden order. In a similar manner, Ref. [129] claims that there are two classes of MBL, “weak” or “strong” depending on whether the spectral functions of l-bits in an MBL system weakly coupled to a bath remain sharp, and that all $d \geq 2$ MBL is weak. We also note that some systems appear to exhibit a mobility edge, an energy (or temperature) above which the system becomes delocalised [93, 130]. In this case, only a portion of the spectrum would be MBL. A similar phenomenon appears with Anderson localisation, but the existence of mobility edges in the many-body case is still debated, with some recent works claiming these do not exist. In particular, Ref. [131] argued that in the presence of a mobility edge, there is a non-zero probability of a thermal excitation sufficiently high-energy to be ergodic, and that such “hot bubbles” might be able to travel through resonant processes, thus delocalising the entire system at any energy. This would not be visible in system sizes accessible to numerics, accounting for why other studies observe mobility edges.

Following on from this, Ref. [39] used similar heuristic arguments to claim that MBL is unstable in $d \geq 2$ to an ergodic inclusion – as well as in $d = 1$ if the localisation length were greater than $1/\ln 2$, or in the presence of interactions which decayed slower than exponentially with distance. The idea is that the ergodic grain would delocalise nearby degrees of freedom, incorporating them into a growing thermal region, which would then have even greater thermalising power due to the increased density of states. However, interactions between the grain and the bulk drop off exponentially with the localisation length. If the density of states grows more quickly than the interactions fall, then the grain will grow forever in a thermal avalanche.

To make this argument more concrete, consider an ergodic grain of linear dimension ℓ_b acting as a bath in an otherwise fully localised system of dimension d . We assume that we have been able to construct perfect l-bits, and that these l-bits have no interactions between themselves; that is, $\varepsilon^{(k)} = 0$ for all $k > 1$ in (2.7). We also assume (for now) that the physical interactions are short-ranged. Therefore, their only couplings are with the bath, and these fall off exponentially with distance,




Figure 2.4: When a rare low-disorder region leads to an ergodic grain in an otherwise localised system – here, modelled by a chain with separate clean and disordered regions – two possible scenarios can result. If the rare region, acting like a thermal bath to the rest of the system, is only able to destabilise a few nearby l-bits, the MBL phase remains stable to the inclusion (left). Far away l-bits experience exponentially small interactions with the bath, and so the bath’s influence spreads logarithmically slowly. On the other hand, the bath may incorporate nearby degrees of freedom (right), which grow the bath’s density of states and therefore increase its thermalising power. As a result, a thermal avalanche spreads through the system and destabilises the MBL phase. Which scenario occurs depends on the dimensionality and the localisation length. [From Ref. [132]. Reproduced with permission from Springer Nature.]

on the scale of the localisation length ξ . The bath is also assumed to be described by a random matrix. Following Fermi’s Golden Rule, if the coupling between an l-bit and the bath exceeds the level spacing in the bath, then the bath will thermalise that l-bit and grow. We can then calculate whether this expansion will eventually halt. After growing by a linear distance ℓ_c , the bath will acquire on the order of $N_c = (\ell_b + \ell_c)^d - \ell_b^d$ new degrees of freedom, increasing the density of states in the bath by a factor 2^{N_c} . At the same time, one would expect the interactions with the closest l-bits to be of a strength $\exp(-\ell_c/\xi)$ relative to the original interaction strength. The growth of the bath will then eventually halt if $2^{N_c} e^{-\ell_c/\xi} \leq 1$ for some finite ℓ_c . Comparing these two factors, it is clear that in $d > 1$ the density of states grows *superexponentially* in ℓ_c and outcompetes the falling interaction strength. But in $d = 1$, then the condition is satisfied if $\xi < 1/\ln 2$, meaning that the bath will only thermalise a finite boundary layer, otherwise the localised phase is unstable. Likewise, if the physical interactions are long-ranged – decaying more slowly than an exponential – then the growing density of states will outcompete the declining interactions for any d . These two situations are illustrated in Fig. 2.4.

This threatens the existence of the MBL phase, because in any sufficiently large system, we will get rare regions with weak disorder arising through pure chance, known as Griffiths regions [133]. These regions will be locally thermal – and act exactly as

the ergodic grain above. This suggests that the localisation-ergodicity transition in the thermodynamic limit may occur at a much higher critical disorder than previously thought, with a sudden breakdown at critical localisation length $\xi_c = 1/\ln 2$. Because one would need very large system sizes to see one of these rare regions, especially when compared to the length scales accessible for exact diagonalisation (~ 20 spins), these effects had not previously been explored in numerics. Additionally, it is argued that (despite the name) the avalanche process would be extremely slow, with the local thermalisation timescale diverging on approach to the avalanche-driven transition from the thermal side, putting it well beyond the times accessible to experiments [37, 134].

Meanwhile, Ref. [135] showed that J.Z. Imbrie’s proof for MBL in $d = 1$ [36] does not generalise to higher dimensions, because resonances – kept controlled and well-separated in 1D – would instead percolate through the system, with arguments similar to those for the expansion of a thermal avalanche. These assorted phenomena – particularly that of the thermal avalanche – have combined to sow significant doubt in the community regarding the existence of MBL in the thermodynamic limit, with active debate still ongoing today. The argument of Ref. [39] makes an assumption that the bath can always be described by a random matrix – but while the matrix has 2^{ℓ_b} elements, it is representing a subsystem described by $O(\ell_b)$ Hamiltonian parameters, and so must contain correlations. Even if the small thermal inclusion is initially random-matrix, this may change as the bath absorbs more spins in the higher-disorder periphery [37]. Despite this, both numerics and experiments seem to show strong evidence of ergodic inclusions resulting in thermal avalanches [38, 42, 43, 45, 132, 136, 137].

As an example, in one of these studies, Ref. [132] considers an effective 1D Bose-Hubbard model constructed in an optical lattice with a site-resolved tunable disorder potential,

$$H_{\text{BH}} = -J \sum_i \left(a_i a_{i+1}^\dagger + \text{h.c.} \right) + \frac{U}{2} \sum_i n_i (n_i - 1) + \sum_i h_i n_i, \quad (2.8)$$

where $a_i^\dagger$ (a_i) is the bosonic creation (annihilation) operator at site i . The first L_{clean} sites have $h_i = 0$, while the remaining L_{dis} have quasiperiodic disorder $h_i =$




Figure 2.5: A 1D Bose-Hubbard model with clean and disordered regions models the effect of a rare ergodic grain on an MBL system. **(a)** Density correlations between all pairs of sites in the system $\langle n_i n_j \rangle_c$, at three different times; the cuts then show this quantity averaged over the clean sites, $g^{(2)}(i)$. At short times, each section exhibits separate dynamics, but later, anticorrelations between the two regions indicate transport across their interface. These decay exponentially into the disordered region, with a decay length ξ_d . **(b)** Correlations within the clean section, with error bars indicating variation between runs. After a few tunnelling times τ , this settles to a steady state, indicating the emergence of a bath. **(c)** When the bath is small ($L_{\text{clean}} = 2$), the correlation decay length ξ_d grows logarithmically in time, but when the bath is larger ($L_{\text{clean}} = 6$), growth accelerates, indicating a thermal avalanche. **(d)** Single-site particle-number entropy growth for three sites at varying distances from the bath. After initially saturating, the entropy begins growing again after a time exponentially large in the distance from the bath – but only when $L_{\text{clean}} = 6$. Note that the three traces are offset for visual clarity. In all figures, solid lines and shaded regions indicate predictions from exact numerics and their uncertainties, respectively. [From Ref. [132]. Reproduced with permission from Springer Nature.]

$W \cos(2\pi\beta i + \phi)$, with $\beta = (1 + \sqrt{5})/2$, to ensure there are no accidental rare regions. This models an ergodic rare region within an otherwise disordered system. The system is initialised in a state with unity filling at each site, and then proceeds to evolve under the Hamiltonian (2.8), after which fluorescence imaging is used to read out site-resolved occupation numbers. In Fig. 2.5(a), the authors calculate the two-site density correlation, $\langle n_i n_j \rangle_c$, and then average this over all clean sites $j \in L_{\text{clean}}$. At short times, the two halves of the system develop their own separate dynamics; however, at long times (compared to the microscopic tunnelling length $\tau = 1/J$), anticorrelations between disordered sites and clean sites indicate particle transport across the interface between the two regions. Matching the picture outlined above, the influence of the thermal inclusion falls off approximately exponentially with distance, from which a decay length ξ_d may be extracted. Fig. 2.5(b) shows that the correlations within the clean region rapidly approach a steady state, confirming that it does indeed act as a bath. Fig. 2.5(c) then depicts the growth of the measured

decay length ξ_d with time, for a small bath ($L_{\text{clean}} = 2$) and a large one ($L_{\text{clean}} = 6$). For the small bath, ξ_d grows logarithmically in time, confirming again the picture of a bath which is only able to thermalise nearby l-bits at an exponentially slow rate. However, a larger bath causes the growth of ξ_d to accelerate at later times – a strong indication that the bath is successfully thermalising adjacent degrees of freedom, and growing in strength. Finally, Fig. 2.5(d) shows the entropy of the particle number distributions on three sites at varying distances from the clean section – while these initially saturate, in the case of a large bath ($L_{\text{clean}} = 6$) the entropy eventually begins growing again, and at a time which grows exponentially with distance from the bath, showing that the bath is successively thermalising sites in the disordered region.

Ref. [132] also found substantial multipoint correlations $\langle n_i n_j n_k \rangle_c$ between the bath and the disordered sites, indicating the presence of non-perturbative many-body processes believed to drive the growth of a thermal avalanche. This highlights the importance of many-body resonances in the breakdown of the localised phase. Ref. [33] considers the many-body Thouless parameter, $\mathcal{G} = \ln \left| \frac{\langle n+1 | \hat{O} | n \rangle}{E_{n+1} - E_n} \right|$, for some local operator $\hat{O}$. In the thermal phase, $\hat{O}$ strong mixes nearby eigenstates, such that $\langle \mathcal{G} \rangle \sim L$, while in the MBL phase, matrix elements are exponentially suppressed such that $\langle \mathcal{G} \rangle \sim -L$. The sign of $\frac{d\langle \mathcal{G} \rangle}{dL}$ therefore provides a sensitive diagnostic of localisation, and the authors find evidence of an intervening thermal phase between two MBL phases with different topological orders even with only very weak interactions (although it is not clear if the localised phase is unstable even to arbitrarily weak interactions at the critical point).

Some works go even further and claim that MBL, as a stable phase of matter in the thermodynamic limit, cannot result from random quenched disorder or at least that the critical disorder is very large [40, 42, 43, 138]. Despite this, study of MBL still remains relevant. Quasiperiodic potentials do not contain rare regions and so are not believed to be susceptible to avalanche instabilities in the same manner, but still exhibit MBL [139–141]. For finite-size systems – or, equivalently, at finite times – we may also be able to observe a prethermal MBL *regime* at disorder strengths below the critical point, which exhibits exponentially slow thermalisation and apparent frozen dynamics [44, 138, 142]. At intermediate disorder strengths, the

phenomenology of the MBL regime is controlled by the proliferation or absence of many-body resonances between eigenstates close together in energy [34, 43, 143–145]. Other works argue that finite-size scaling has been incorrectly applied and conflicts with analytical results proving the existence of MBL [146]. It is still a substantial open question, therefore, whether MBL indeed survives in the thermodynamic limit, and if it does, how to calculate the true critical disorder strength. Furthermore, even in finite size systems and at finite times the behaviour of the MBL regime is still unclear.

2.3 Quantum Scars

The discovery of MBL, a robust phase of matter in which every eigenstate violates the ETH and thermalisation is seemingly prevented entirely, naturally led to the question of whether there exist other classes of systems which violate the ETH. In particular, could ETH-violating states coexist with thermal states in a physically reasonable non-integrable system? Such a system would sit in between the regimes of integrability and quantum chaos, breaking the strong ETH – which requires the ansatz to hold for *every* state with finite energy density – but not the weak ETH, which only requires it for *almost all* such states. Indeed, Ref. [6] stated that, “it would indeed be very interesting if there existed some experimental isolated systems that could be shown to fail to thermalise from certain states that were carefully prepared.”

The experimental discovery in recent years of so-called quantum many-body scars in a Rydberg atom simulator [9] showed that this phenomenon, known as weak ergodicity breaking, is indeed possible. The system, which is equivalent at low energies to a spin-1/2 chain in which excitations are forbidden from existing on adjacent sites, demonstrated atypical dynamics with oscillatory behaviour and periodic revivals from charge density wave (CDW) initial states, but thermalising behaviour when starting in other states; see Fig. 2.6 and Fig. 2.7. These states were dubbed quantum many-body scars (QMBS), due to their similarity to exceptional states in single-particle quantum billiards. This stands in contrast to MBL, integrable, and ergodic systems, where the qualitative features of the dynamics do not depend on initial conditions: the variable nature of dynamics is a key signature of quantum scarring, and this behaviour even holds implications for the general validity of




Figure 2.6: (a) Experiments with a Rydberg atom simulator in the Rydberg blockade regime, in which excitations are forbidden from existing on adjacent sites, show oscillations between two charge density wave states. (b) Domain wall density after a quench, showing slow decay of oscillations. (d) Numerical simulations of 25 Rydberg atoms, showing also slow growth in entanglement entropy; taking into account the full $1/R^6$ interactions causes the oscillations to decay more quickly than assuming a perfect blockade. [From Ref. [9]. Reproduced with permission from Springer Nature.]




Figure 2.7: Numerical simulations of the PXP model, a spin-1/2 chain designed to capture the key features of the Rydberg blockade regime. The entanglement entropy and spin-spin correlation show clear oscillations, but *only* from certain charge density wave initial states $|Z_{2,3}\rangle$. These states also show slow growth of entanglement. Here, ΔS is the rate of entanglement entropy growth. [From Ref. [18]. Reproduced with permission from Springer Nature.]

statistical mechanics [6].

In subsequent sections we will look at the phenomenology of quantum scars, the proposed theoretical underpinnings, and the connections between QMBS and their single-particle equivalents.

2.3.1 Persistent oscillations in chaotic systems

An atom in a Rydberg state has one or more electrons in a highly-excited orbital (that is, with large principal quantum number n). Rydberg states have a very long lifetime and, due to the large radius of high- n electron orbitals, result in strong repulsive van der Waals interactions. Ref. [9] realises a chain of atoms which are resonantly driven between the ground state and a Rydberg state, forming effective spin-1/2 particles; writing the ground state as $|\circ\rangle$ and Rydberg state as $|\bullet\rangle$, the Hamiltonian may be written,

$$H_{\text{Ryd}} = \sum_j \frac{\Omega}{2} \sigma_j^x - \sum_j \Delta n_j + \sum_{j < k} V_{jk} n_j n_k, \quad (2.9)$$

where $\sigma^x = |\circ\rangle\langle\bullet| + |\bullet\rangle\langle\circ|$, $n_j = |\bullet\rangle\langle\bullet|$, and $V_{jk} = C|j - k|^{-6}$ with $C > 0$. When C is large, the interaction term introduces an extreme energy cost to adjacent excitations, a regime known as the Rydberg blockade.

Fig. 2.6(a) shows how remarkably, when initialised in the charge density wave state $|\mathbb{Z}_2\rangle = |\circ\bullet\circ\bullet\dots\rangle$ and placed in the Rydberg blockage regime with $\Delta = 0$, the system exhibited coherent and persistent oscillations between $|\mathbb{Z}_2\rangle$ and the translated state $|\mathbb{Z}'_2\rangle = |\bullet\circ\bullet\circ\dots\rangle$, yet in any other state it appeared to rapidly thermalise. Likewise, as shown in Fig 2.6(b), the domain wall density following the quench showed continuing oscillations, although these decayed slowly over time. Numerical simulations, in Fig 2.6(d), also showed slow growth of entanglement, and that the oscillations were enhanced when the more realistic $1/R^6$ interactions were replaced by a strong nearest-neighbour interaction which entirely forbade adjacent excitations. As explained above, these results were extremely surprising: until this point, all many-body systems had been found to exhibit either chaotic or integrable dynamics at any one energy, but in this simulator, two states at the same energy could have wildly different behaviours.

In the Rydberg blockade regime, the system can be well approximated by the




Figure 2.8: (a) Bipartite half-chain entanglement entropy and (b) overlap with the charge density wave state $|\mathbb{Z}_2\rangle$ against energy for all eigenstates in the PXP model. The labelled band of states, exhibiting low entanglement entropy, near-equal energy spacing, and atypically high overlaps, explain the persistent oscillations in Fig. 2.7. [Reprinted figure with permission from Ref. [19]. Copyright (2018) by the American Physical Society.]

following Hamiltonian, known as the PXP model [18, 19]:

$$H_{\text{PXP}} = \sum_j P_{j-1} \sigma_j^x P_{j+1}, \quad P_j = \frac{1 - \sigma_j^z}{2}. \quad (2.10)$$

This model exhibits tightly constrained dynamics, as a spin cannot flip unless both of its neighbours are in the ground state. Crucially, as seen in Fig 2.7, it captures the key features of the Rydberg blockade regime, including the unusual dynamics. Numerical simulations of this model following a quench from the state $|\mathbb{Z}_2\rangle$ also show the rate of entanglement growth oscillating at the same frequency as the domain wall density, and furthermore persistent (but weaker) oscillations starting from the state $|\mathbb{Z}_3\rangle$, in which every third atom is excited. However, the period-4 charge density wave state $|\mathbb{Z}_4\rangle$ and other randomly chosen states are seen to rapidly thermalise. Exact diagonalisation of the PXP model shows a band of ETH-violating eigenstates [18, 19] sitting in the middle of an otherwise chaotic system, as seen in Fig. 2.8. These states have low entanglement entropies and unusually high overlaps with $|\mathbb{Z}_2\rangle$, and are also spaced nearly equally in energy, which explains the oscillations – a superposition of

states with an energy spacing Δ will return to the original state after a period $1/\Delta$, an occurrence known as a many-body revival. These states were dubbed many-body quantum scars, in analogy with “scarred” states in quantum billiard models that are the remnants of unstable but periodic classical orbits [20, 147–150] – a connection I will elaborate on in Sec 2.3.3.

The PXP model has now become the paradigmatic example of QMBS, with oscillatory or long-lived dynamics in an otherwise chaotic system starting from certain special experimentally-preparable states, and this behaviour explained by a set of exceptional energy eigenstates which violate the ETH. Another model that has attracted considerable attention is the AKLT model, itself already a paradigmatic toy model of symmetry-protected topological (SPT) order (see Sec. 2.4.2), which was recently found to contain a “tower of states” related by a raising/lowering operator algebra [10]. Unlike the states in the PXP model, the AKLT scars are perfect – exact analytical expressions are known for them, they have integer energy spacing (in the appropriate units), and they provably survive to the thermodynamic limit, at least at the finely-tuned fixed point of the AKLT model. (Their survival in the wider AKLT phase is still an open question.) Notably, the model has continuous $SU(2)$ symmetry, which is known to forbid MBL [151, 152] – but this does not seem to have prevented weak ergodicity breaking.

Just as Rydberg atom simulators led to the first discovery of QMBS, experiments continue to provide new insights into this novel form of ergodicity breaking. For example, [153] used a superconducting quantum processor (depicted in Fig. 2.9),




Figure 2.9: A superconducting processor containing 36 qubits in a 6×6 grid (grey) and 60 couplers (yellow), which models a chain inspired by the Su–Schrieffer–Heeger (SSH) model of polyacetylene. The chain has alternating strong and weak couplings J_a (dark yellow) and J_e (light yellow), forming oscillating dimers in the absence of random cross-couplings J_x (blue). [From Ref. [153]. Reproduced with permission from Springer Nature.]

with 36 qubits in a 6×6 grid and 60 couplers (which are two-level systems, linking adjacent qubits), to simulate a chain inspired by the Su–Schrieffer–Heeger (SSH) model of polyacetylene. This chain has alternating strong and weak couplings, as well as random cross-couplings between well-separated sites which break the integrability, and it is best to think about this chain in terms of the dimers connected by the strong couplings, which can be in the states $|d_0\rangle = |00\rangle$, $|d_1\rangle = |11\rangle$, $|d_-\rangle = |01\rangle$, and $|d_+\rangle = |10\rangle$. In the absence of the cross couplings, the chain stably oscillates between two states $|\Pi\rangle = |d_-d_+d_-d_+\dots\rangle$ and $|\Pi'\rangle = |d_+d_-d_+d_-\dots\rangle$. In this limit, the system is fully integrable, but remarkably, this particular oscillation persists even as the integrability-breaking perturbations are included, despite the system thermalising from any other initial state. Fig. 2.10 shows the dynamics of the superconducting processor following a quench to either $|\Pi\rangle$ or an arbitrary spin product state with unity filling, showing clearly that significant oscillations only result from $|\Pi\rangle$ or its counterpart. We additionally note that the PXP model has been realised on a tilted optical lattice [154], where a linear potential creates an emergent kinetic constraint – this should be contrasted with the Rydberg atom simulators of Ref. [9], which have strong $1/R^6$ van der Waals interactions that are then *approximated* by the PXP constraint.

2.3.2 Theories of quantum scarring

While there exist explanations for certain classes of quantum scars, there is still not a unified picture for quantum scarring as a whole, although it has now been shown that many-body revivals can only be explained by a band of exceptional scar states [155]. In many cases, just like in the AKLT model, scars take the form of a tower of states with constant energy separation [11–16]. However, this is not the only paradigm: quantum scars have also resulted from the existence of decoupled or weakly coupled Krylov subspaces [156]. This is a phenomenon in which the space of all states obtained by repeated action of the Hamiltonian on some state $|\psi\rangle$ – that is, $\text{span}\{|\psi\rangle, \mathcal{H}|\psi\rangle, \mathcal{H}^2|\psi\rangle, \dots\}$ – does not span the Hilbert space, which means that the dynamics starting from $|\psi\rangle$ remain within this subspace for all time. Such subspaces may be integrable or chaotic, coexisting within the same system, and hence some may be classified as a set of quantum scars. In the extreme case, this results in Hilbert space fragmentation [157, 158], in which the Hilbert space




Figure 2.10: Oscillating quantum scars in the superconducting quantum processor from Fig. 2.10. **(a, b)** Dynamics of the on-site occupation number n_i , starting respectively from either the scarred state $|\Pi\rangle$ or an arbitrary spin product state with unity filling, with $L = 30$. In (a), oscillations between $|\Pi\rangle$ and $|\Pi'\rangle$ are clearly observed, but in (b) the system appears to decay to thermal equilibrium. **(c, d)** Generalised imbalance $I(t) \propto (1/L) \sum_i \langle \sigma_i^z(t) \rangle \langle \sigma_i^z(0) \rangle$, for the two states in (a) and (b) respectively. For the scarred state in (c), the imbalance shows clear periodic oscillations, but for the random initial state in (d), the imbalance rapidly decays to zero. **(e, f)** Fourier transform of the imbalance $I(t)$, $g_\alpha(\omega)$, for the two states respectively. Both traces exhibit peaks at a particular frequency ω_1 , but the peak is much larger for the scarred state in (e). **(g)** The height of the Fourier peak $g_\alpha(\omega_1)$ for the scarred state, extracted from (e), against the ratio between intra- and inter-dimer couplings J_a/J_e . **(h)** Height of the Fourier peak, $g_\alpha^2(\omega_1)$, for all states with unity filling. Green stars highlight the two quantum scars $|\Pi\rangle$ and $|\Pi'\rangle$, which are the only states with a substantial peak. In the insets of (c-f) and in (g), solid lines depict exact numerics for a smaller system $L = 20$. [From Ref. [153]. Reproduced with permission from Springer Nature.]

is split into exponentially many disconnected subspaces (some of which may only contain one or a handful of states). However, such scars are not usually equally spaced in energy and do not always lead to revivals. Moreover, it is also possible to systematically construct models that violate the ETH, using a technique known as projector embedding [17] to insert arbitrary subspaces into otherwise chaotic systems. While of theoretical interest, this final construction provides little insight into the physical origin of scars.

It has been shown that many of the aforementioned towers of states may in fact derive from broken symmetries in the Hamiltonian [15], due to the presence of spectrum-generating algebras (SGAs). Specifically, consider a Hamiltonian H_{sym} with some continuous non-Abelian symmetry G : take $G = \text{SU}(2)$, which has raising and

lowering operators $Q^\pm$ and a generator Q^z such that $[Q^z, Q^\pm] = \pm Q^\pm$. In this case, the spectrum will contain degenerate multiplets of eigenstates, with each multiplet labelled by the Casimir operator Q^2 and the states within each multiplet by Q^z . By applying an additional term H_{SG} e.g. ωQ^z , the degeneracy may be broken without altering the eigenstates, creating towers of states. Finally, some integrability-breaking term H_A is chosen which annihilates one particular tower, leaving it unchanged as a set of equally spaced quantum scars in an otherwise chaotic spectrum. This is a generalisation of previous work which connected scars to an effective $\text{SU}(2)$ picture [14, 159]; the difference here is that this construction can be applied to any continuous non-Abelian symmetry.

Despite being the first scarred system to be discovered experimentally, the PXP model on the other hand so far appears to have defied neat categorisation, and the origin of the scars within it is still an open question. Unlike in the SGA construction above, the scars in it are not known exactly, and have only been approximated analytically. In spite of this, some progress has been made towards understanding them. It has been shown that the PXP model may be deformed to an integrable point, but such a deformation destroys the scars, implying this is not connected to their presence [19]; but other deformations appear to move the model closer to integrability *and* enhance the oscillations [160]. The scars also appear to be linked to critical states at an Ising-type phase transition induced by a transverse field [161] (although these critical states exhibit thermalising dynamics), while it has been shown that PXP scars can survive a quench through a (different) critical point [162]. Additionally, a certain weak quasi-local deformation to the model leads to emergent $\text{SU}(2)$ -spin dynamics, making the revivals virtually perfect [163] – which could mean that the scars do in fact originate from an SGA. The PXP model also continues to exhibit scarring when looking at higher-spin equivalents; it has been suggested recently that combining dynamical constraints with large local degrees of freedom makes systems amenable to quantum scarring [164].

Finally, let us talk about the interplay between disorder, MBL, and quantum scars. This is an important question for experiments, as certain quantum simulators which may be used to investigate scars exhibit natural variations due to imperfections. It has been shown that scars may exist in disordered systems [13, 165], but these

works concerned models in which the disorder was fine-tuned to not affect the scars. Recent work on the PXP model has shown that not only can the scars survive weak generic disorder, but that such disorder helps uncover a deeper multi-tower structure which had previously gone unnoticed [166]. Increasing the disorder strength then leads to fully ergodic behaviour (as the scars are destroyed), followed by constrained and then unconstrained MBL. However, there is still some dispute over whether such constrained systems are able to localise at all. While Ref. [167] claims that transverse field disorder in the PXP model leads to a stable MBL regime, it also notes that the constraints can frustrate the ability of spins to align with this field, and instead lead to ergodic or critical behaviour. Likewise, Ref. [168] finds that the combination of constraints and disorder in fact leads to an induced interaction, which can prevent MBL entirely in the thermodynamic limit no matter how strong the disorder.

As quantum scarring is a relatively young field, numerical techniques for studying it remain limited: most studies have either shown the existence of scars analytically, or used exact diagonalisation or certain MPS-based techniques. One recent helpful development is a variant of DMRG, dubbed DMRG-S, which is able to systematically locate low-entanglement states and derive an MPS representation [169]. One other technique, the time-dependent variational principle (TDVP) [170, 171], allows one to construct effective equations of motion within a variational manifold, which has led to new theoretical insights via connections to QMBS's single-particle namesake – we will discuss these in the next section.

2.3.3 Connection to single-particle quantum scars

Consider a classical particle moving around inside a box, bouncing elastically off the walls with equal angles of incidence and reflection – a problem known as dynamical billiards. Depending on the shape of the box, the dynamics could be integrable or chaotic. On one hand, a circular box results in a predictable trajectory (e.g. every bounce has the same angle of incidence) and so the ball does not fully explore the phase space⁴ available to it, entirely avoiding the centre of the circle in almost all cases. On the other hand, by extending this circle into a lozenge made of a rectangle with semicircular ends – known as the Bunimovich stadium [172] – the dynamics becomes completely unpredictable; see Fig 2.11(a). Particles launched close together

⁴The phase space here is the set of all points in the box combined with the direction the ball moves in.




Figure 2.11: (a-b) The motion of a particle launched along an arbitrary trajectory in the Bunimovich stadium (a) is unpredictable. However, there exist certain periodic orbits (b, red) which make up a vanishing fraction of phase space. These orbits are unstable: a slight deviation (b, blue) will eventually result in chaotic dynamics. (c-d) A typical eigenstate of the quantum stadium (c) appears random. However, there exist certain rare eigenstates (d) clustered around the unstable periodic orbits of the classical stadium, which make up a vanishing fraction of the spectrum as $\hbar \rightarrow 0$. These are identified as quantum scars. (e) Time-evolved fidelity of a wavepacket launched either at a 45 deg angle (red) or vertically between the straight walls (blue). The former rapidly decays, while the latter shows periodic revivals. (f) Spectral decomposition of the two initial states from (e). The state which leads to revivals (blue) exhibits strong overlaps with equally-spaced eigenstates, which are exactly the quantum scars of (d). [From Ref. [20]. Reproduced with permission from Springer Nature.]

will diverge exponentially and explore the entirety of phase space. Likewise, when one quantises these problems, turning them into the case of a single particle in an infinite well, the systems match their classical counterparts. There is no ETH or a notion of thermalisation for single particle systems, but in chaotic cases such as the Bunimovich stadium, the eigenstates display level repulsion and appear essentially random [20] – see Fig 2.11(c) – while the circular well will have analytic solutions with well-defined quantum numbers.

Yet if we look closely at the Bunimovich stadium in Fig 2.11(b), we notice that we can bounce a particle vertically between the two straight walls of the rectangle to obtain a perfectly periodic orbit, albeit an unstable one: it is not perfectly ergodic. These orbits have vanishing measure in phase space, so the stadium is still considered chaotic. Remarkably, when we look at the energy eigenstates of the stadium, we find certain special states whose probability density clusters around the unstable orbits [20]; one of these is shown in Fig 2.11(d). Like the classical orbits, these states are a vanishing proportion of the full spectrum, disappearing in the classical limit

$\hbar \rightarrow 0$ (equivalent to going to very high energies). Furthermore, if we launch a tight wavepacket along one of these orbits, we observe quantum revivals in Fig 2.11(e), and the eigenstate decomposition of this initial state shows a strong overlap with the aforementioned special states in Fig 2.11(f). This was termed quantum scarring, and is the namesake of QMBS due to the strong similarity between the two phenomena. But there is one piece missing: single-particle quantum scars corresponded to unstable periodic orbits in classical systems, but there are no obvious classical counterparts to systems like the PXP or AKLT model, which are quantum in origin.

There is a way to create one, though. The time-dependent variational principle (TDVP) can be used to determine the best approximation to full Hilbert space dynamics in a state constrained to a variational manifold [170, 171] such as an MPS of fixed bond dimension. TDVP yields a set of (coupled) ordinary differential equations for the parameters of the state, which then may be integrated to obtain a semiclassical trajectory through the manifold. It also naturally respects conserved quantities, providing a symplectic integrator is used to obtain the final trajectory. By using a simple ansatz for the PXP model, expressed in the form of an MPS, Ref. [173] was able to apply the TDVP to the model and uncover *unstable, periodic orbits* passing through the same two $|\mathbb{Z}_2\rangle$ states between which oscillations were observed in Rydberg simulators [9, 18], thus establishing the missing link between the single- and many-body cases. One can also directly quantify the error between the variational approximation and full many-body time evolution – this was shown to be substantially reduced by the deformations applied in Ref. [160]. By converting quantum mechanical time evolution into classical equations of motion, within which periodic orbits can be clearly identified and analytically studied, TDVP has been useful for investigating the changing behaviour of QMBS across a wide variety of parameter regimes [162] as well as systematically constructing new families of scarred Hamiltonians [174]. Furthermore, Ref. [175] uncovers a set of quasimodes in the PXP model which also arise from a “requantisation” of the semiclassical TDVP system, demonstrating a correspondence between classical and quantum cases similar to that observed in single-particle scars.

2.4 Topological Order

Many-body localisation and quantum many-body scars both concern highly excited states, up to and including those at infinite temperature. However, we now look at a very different topic: that of topological order, which captures how systems can undergo phase transitions at *zero* temperature without spontaneous symmetry-breaking or a local order parameter. Topologically non-trivial systems have a gapped ground state, which may be degenerate, and which cannot be continuously deformed to a trivial product state by perturbations to the Hamiltonian without closing the gap. The gap and the degeneracy as well as certain signatures of this order are said to be topologically protected, and the closure of the gap is identified with a *topological* phase transition. In this thesis, we are more interested in symmetry protected topological (SPT) order, in which only perturbations which respect a particular symmetry of the Hamiltonian will preserve topological order [176]. SPT order is typically associated with protected edge modes, degenerate entanglement spectra, and non-local order such as string order or topological entanglement [21–27], but at finite temperature interactions with the bulk cause the edge modes to decohere, and other signatures of SPT order are lost. Unlike l-bits, which can only store *classical* information efficiently, edge modes can store quantum information at zero temperature. However, if these could be preserved beyond the ground state, they could be used to build a robust quantum memory [124]. Quenched disorder has been shown to preserve signatures of SPT order even up to infinite temperature [29–31, 35, 177, 178], while topological quantum scars have been inserted into the spectrum of otherwise chaotic models [179, 180]. The interplay between SPT order and ergodicity breaking phenomena is a core thread in this thesis: in the following I will briefly review the main properties of SPT order, as well as methods for stabilising it at finite temperature.

2.4.1 Classification of topological order

Until the discovery of topological order, phase transitions were believed to be fully explained by Landau’s paradigm of spontaneous symmetry breaking, in which all phases could be classified by a local order parameter. However, beginning with the discovery of the fractional quantum Hall effect in the 1980s, it was becoming clear that there was more to quantum mechanical phases than symmetry breaking, and

that an entirely new theory of phase transitions would be needed to describe such phenomena. This led to the idea of topological order, which describes the phases of gapped quantum systems at zero temperature and classifies their ground states. A perturbation to such a system, so long as it does not close the gap, will only result in a local restructuring of the ground state, and so all systems that can be smoothly transformed into one another without closing the ground state gap form an equivalence class, labelled by topological invariants. This is reminiscent of topological invariants in geometry, for example the number of holes in a manifold, which cannot be altered by a smooth deformation. In particular, the topological order of a state cannot be changed by finite-time evolution under a local Hamiltonian or, equivalently, under a finite-depth circuit of local unitaries.

In one dimension, the ground state of gapped Hamiltonians, which follow an area law and have short-range correlations, can be described by a matrix product state (MPS),

$$|\psi_{\text{MPS}}\rangle = \sum_{\{j\}} \text{Tr} \left(A_{j_1}^{[1]} A_{j_2}^{[2]} \dots A_{j_L}^{[L]} \right) |j_1 j_2 \dots j_L\rangle \quad (2.11)$$

where for each site ℓ , $A_{j_k}^{[\ell]}$ are a set of d tensors $A_1^{[\ell]}, A_2^{[\ell]}, \dots, A_d^{[\ell]}$, each of size $\chi \times \chi$, where χ is known as the bond dimension and d is the local Hilbert dimension. (The above equation is for periodic boundary conditions; for open boundaries, the matrices on the ends are replaced by vectors and there is no trace.) By applying a renormalisation group technique to these states, it is possible to show that there is no topological order in one dimension: all possible ground states flow towards the same fixed point, in the equivalence class of a product state, and are therefore in the trivial phase [176]. However, this all changes if we introduce symmetries. If we demand that the Hamiltonian respects a particular symmetry, then we likewise have to respect that symmetry during the RG, and the proof is no longer valid. Instead, we can classify phases by the action of this symmetry (or symmetries) on the state; this may be combined with spontaneous symmetry breaking, in which case we look at the action of the unbroken subgroup. These phases are called symmetry-protected topological (SPT) phases, and are only robust to perturbations that respect the protecting symmetry.

If we act on an MPS tensor $A_{j,\alpha\beta}$ with a symmetry $g \in G$ that acts on-site (such

as $\mathbb{Z}_2$ spin flip symmetry), we can “push” it through the tensor onto the bonds:

$$u(g)_{jk} A_{k,\alpha\beta} = a(g) R(g)_{\alpha\alpha'}^{-1} A_{j,\alpha'\beta'} R_{\beta'\beta}(g) , \quad (2.12)$$

where $u(g)$ is a representation of g on the local Hilbert space, $a(g)$ a one-dimensional representation (i.e. a complex phase), and $R(g)$ a representation on the bonds. But because a quantum state is only defined up to an overall phase, the $R(g)$ can form a *projective* representation, such that $R(g_1)R(g_2) = \omega(g_1, g_2)R(g_1g_2)$. The equivalence classes of projective representations are labelled by elements of the second cohomology group $\omega \in H^2(G, \mathbb{C})$, and it is these that define the SPT phase. Using these techniques, Refs. [176, 181] exhaustively classify SPT phases in 1D spin systems (and by the Jordan-Wigner transform, fermionic systems too).

At the same time, other groups made similar insights into SPT order, through from slightly different perspectives. Refs. [24, 182] instead classified SPT phases by their entanglement spectra (2.4), again using projective representations of symmetries. Meanwhile, Refs. [183, 184] looked at the edge modes associated with topological order in chains of Majorana fermions with open boundary conditions, and the action of symmetries on these edges, finding that there were just 8 distinct SPT phases in fermionic chains when interactions were present.

2.4.2 Edge modes, entanglement spectra, and string order

SPT order typically leads to phenomena such as protected edge modes, degeneracies in the entanglement spectrum, and non-local ordering such as string order or topological entanglement [21–27]. Consider, for example, the celebrated AKLT (Affleck, Kennedy, Lieb, Tasaki) model, a spin-1 chain with the following Hamiltonian:

$$H_{\text{AKLT}} = \sum_j P_{j,j+1}^{(2)} , \quad P_{j,j+1}^{(2)} \propto \vec{S}_j \cdot \vec{S}_{j+1} + \frac{1}{3} \left(\vec{S}_j \cdot \vec{S}_{j+1} \right)^2 . \quad (2.13)$$

Here, $P_{j,j+1}^{(2)}$ is the projector from two spin-1s onto spin-2. Since a projector is positive semi-definite, the ground state is the state annihilated by all projectors. This can be satisfied by splitting each spin-1 into two virtual spin-1/2 degrees of freedom, and then placing virtual spins from adjacent sites into spin-singlet states. As each singlet has total spin 0, and every adjacent pair of physical spin-1 contains one of these

singlets, it is impossible for the total spin to be 2. Therefore by construction, this state is annihilated by H_{AKLT} and is the ground state. It can also be represented exactly by an MPS of bond dimension $\chi = 2$, given by $A^+ = \sqrt{\frac{2}{3}}\sigma^+$, $A^0 = -\sqrt{\frac{1}{3}}\sigma^z$, $A^- = -\sqrt{\frac{2}{3}}\sigma^-$, and analysis of this state shows that it is in an SPT phase [21, 22, 24] now known as the Haldane phase, protected by any one of three symmetries – the $\mathbb{Z}_2 \times \mathbb{Z}_2$ group of π -rotations around the x , y , and z axes; time reversal symmetry; and spatial inversion [25].

However, now consider what happens when we introduce open boundary conditions. Every virtual spin-1/2 in the bulk is paired into a singlet, but at the edges these are left dangling. Remarkably, spin-1 degrees of freedom have been *fractionalised* into spin-1/2s, whose signatures can be measured in experiment. Crucially, these edge modes are unconstrained by the Hamiltonian, and take zero energy to flip, meaning that the ground state is four-fold degenerate, with the ground state manifold encoding the edge degrees of freedom. If we move away from the AKLT point (for example, by changing the coefficient of the biquadratic term $(\vec{S}_j \cdot \vec{S}_{j+1})^2$), this degeneracy is preserved in the ground state so long as we do not close the gap and leave the SPT phase.

We can also look at measuring the entanglement across a bipartition of the chain. Because we are always cutting a single spin singlet, there is a constant $S = \ln 2$ entanglement entropy, and furthermore the entanglement spectrum is exactly two-fold degenerate. Refs. [24, 182] show that this entanglement degeneracy is more general, and persists even in the absence of edge modes (which only appear in certain cases), allowing one to classify SPT phases through the projective representations of symmetries acting on the MPS representation of ground states (similarly to Refs. [176, 181]). Where the irreducible representations of these symmetries are all multi-dimensional, this leads to protected degeneracies in the entanglement spectrum. This can be illustrated by considering a variation of the MPS ansatz, where the matrices $A_m^{[\ell]}$ are split into on-site matrices $\Gamma_m^{[\ell]}$ and a bond matrix $\Lambda^{[\ell]}$ such that $\Gamma_m^{[\ell]}\Lambda^{[\ell]} = A_m^{[\ell]}$. The eigenvalues of $\Lambda^{[\ell]}$ are exactly the Schmidt coefficients of a bipartition at that bond, and so give the entanglement spectrum. Furthermore, and suppressing the ℓ -dependence, we can always choose Γ_m , Λ such that the

transfer matrix,

$$T_{aa';bb'} = \sum_m \Gamma_{ab}^m (\Gamma_{a'b'}^m)^* \Lambda_b \Lambda_{b'} , \quad (2.14)$$

has a unique largest eigenvalue 1 with right eigenvector $\delta_{bb'} = \mathbb{I}$.

To summarise an example from Ref. [24], consider the action of a bond-centred inversion symmetry, which acts locally as $\Gamma_m \rightarrow \Gamma_m^T = e^{i\theta_I} U_I^\dagger \Gamma_m U_I$ with U_I unitary [c.f. $R(g)$ in Eq. (2.12)]. By iterating this twice, we arrive at $\Gamma_m = e^{2i\theta_I} (U_I U_I^*)^\dagger \Gamma_m U_I U_I^*$, which in fact implies that $U_I U_I^*$ is an eigenvector of the transfer matrix with eigenvalue $e^{2i\theta_I}$. But by assumption, the transfer matrix has a unique unit-modulus eigenvalue 1, with eigenvector $\mathbb{I}$, so $e^{2i\theta_I} = 1$ and $U_I U_I^* = e^{i\phi_I} \mathbb{I}$ for some phase ϕ_I . Since U_I is unitary, then $U_I^T = U_I e^{-i\phi_I}$ and $e^{-2i\phi_I} = 1$, which crucially implies that $\phi_I = 0$ or π . ϕ_I cannot change continuously, and so can only change at a phase transition, hence defining two SPT phases. If $\phi_I = \pi$, U_I is antisymmetric, and since the states in a Schmidt decomposition also transform under U_I , this causes the entanglement spectrum to be doubly degenerate. One can show that inversion symmetry acts upon the AKLT matrices with $\phi_I = \pi$, and therefore this degeneracy is associated with the *entire* Haldane phase.

Since $\sigma^\pm (\sigma^z)^n \sigma^\pm = 0$ for any $n \geq 0$, we can also see from the MPS representation that if we write the ground state out in the spin product-state basis and filter out the sites in the $|0\rangle$ state, every basis state will be formed of alternating $|+\rangle$ s and $|-\rangle$ s (for example, $|+00-0+00-+-\rangle$). This defines a non-local order parameter, known as the string order parameter, which is constructed to be equal to 1 when this condition is satisfied and 0 otherwise. Again, this parameter remains non-zero throughout the entire SPT phase. This can be generalised: Ref. [185] provides a construction for detecting any SPT phase in one dimension using a string order parameter.

These features of the AKLT model are exemplary of the phenomena associated with SPT order in general. For example, the transverse field Ising model (TFIM) can be written in terms of Majorana fermions using the Jordan-Wigner transform, with a pair of fermions making up each spin-1/2:

$$H_{\text{TFIM}} = \sum_j -J \sigma_j^z \sigma_{j+1}^z - \Gamma \sigma_j^x \quad (2.15)$$

In the ferromagnetic phase these fermions become paired up analogously to those in

the AKLT ground state – again leaving an unpaired fermion at each end of an open chain. This results in a two-fold ground state degeneracy, protected by $\mathbb{Z}_2$ spin flip symmetry, and a single spin-1/2 edge zero mode (bilocalised to the two ends of the chain). However, unlike in the AKLT model, where the protected degeneracy only exists in the ground state, in the TFIM it is replicated through the entire spectrum: the edge mode forms what is known as a strong zero mode.

A strong zero mode results from any operator Ψ satisfying the following three properties [128, 186]:

1. $\|[H, \Psi]\| = \mathcal{O}(e^{-\alpha L})$ as $L \rightarrow \infty$ for some $\alpha > 0$ and a chain of length L ,
2. $\{\mathcal{D}, \Psi\} = 0$ for some finite symmetry $\mathcal{D}$ (that is, where $[H, \mathcal{D}] = 0$ and $\mathcal{D}^m = \mathcal{I}$ for some integer $m > 1$),
3. $\Psi^n \propto \mathcal{I}$ for some integer $n > 1$.

Because Ψ anticommutes with the symmetry, its action maps states onto a different symmetry sector, but the first property ensures that this new state will also be an eigenstate (up to exponentially small corrections). The third property simply ensures the new states are normalisable. In the thermodynamic limit, therefore, there is an exact pairing of states across symmetry sectors.

In the TFIM, it is easy to show that in the $\Gamma = 0$ limit, $\Psi = \sigma^z$ satisfies these properties for the spin-flip symmetry $\mathcal{D} = \prod_j \sigma_j^x$, and in fact is an *exact* zero mode with $[H, \Psi] = 0$. One can perturbatively construct the operator in the entire $|J| > |\Gamma|$ phase, solving order by order to eliminate $[H, \Psi]$, as,

$$\Psi \propto \sigma_1^z + \frac{\Gamma}{J} \sigma_1^x \sigma_2^z + \frac{\Gamma^2}{J^2} \sigma_1^x \sigma_2^x \sigma_3^z + \cdots = \sum_{j=1}^L \left(\frac{\Gamma}{J}\right)^{j-1} \left(\prod_{k < j} \sigma_k^x\right) \sigma_j^z. \quad (2.16)$$

This now has nonzero commutator with the Hamiltonian, but still satisfies all three conditions, and remains localised to the edge with an exponentially decaying tail. This means that the edge mode can be addressed at the edge, without disturbing the state deep into the bulk, and moreover that the physical edge spin σ_1^z has an infinite coherence time in the thermodynamic limit at any temperature. It is non-trivial that this should exist – for example, condition 3 requires precise cancellation between a large number of terms – and the construction does not work starting from e.g. σ_2^z (although it will for σ_L^z) or for all models with edge modes.




Figure 2.12: (a) Time evolution of $\langle \sigma_1^z(t) \rangle$ following a quench from a random σ^z basis product state in the $h_j = 0$ and $g > 0.5$ regime, and the corresponding Fourier transform, showing peaks due to single-particle modes. The two sharp peaks near $\omega = \pi$ are attributed to the π edge modes. (b) The single-particle quasienergy spectrum against g for $L = 6, 12$, and 18 . The spectra exhibit a splitting Δ between the edge modes, and a bulk-edge gap Ξ . (c) Sections of the spectra near $\omega = \pi$, for various system sizes L , showing that the edge mode splitting Δ shrinks with increasing L . Plotting Δ against L for various g , we see that the splitting shrinks exponentially with L for g sufficiently far from the critical point, exactly as predicted by theory. [From Ref. [187]. Reprinted with permission from AAAS.]




Figure 2.13: (a) In the Majorana representation, the Ising model has two edge modes, each formed of a sum of single fermions with decreasing weights away from the edges. In the spin representation, these single fermions are replaced by non-local Pauli strings C . (b) The magnitudes of the coefficients in the Pauli string expansion of the edge modes, as a function of how far into the bulk these Pauli strings extend (here labelled “Correlator index, n ”). This shows a clear exponential decay, matching predictions. [From Ref. [187]. Reprinted with permission from AAAS.]

These edge modes have been detected directly in experiment. For example, Ref. [187] implements a kicked Ising model, a Floquet system with time evolution operator,

$$U_F = e^{-\frac{i}{2} \sum_{j=1}^L h_j \sigma_j^z} e^{-\frac{i\pi J}{2} \sum_{j=1}^{L-1} \sigma_j^z \sigma_{j+1}^z} e^{-\frac{i\pi g}{2} \sum_{j=1}^L \sigma_j^x}, \quad (2.17)$$

on an array of superconducting qubits. With $h_j = 0$, this model is integrable with a $\mathbb{Z}_2$ spin-flip symmetry, and exhibits two different phases: with $g > J = 0.5$, each eigenstate of U_F with quasienergy θ (such that the eigenvalue is $e^{i\theta}$) has a partner with quasienergy $\theta + \pi$, connected by a pair of Majorana edge modes (MEMs) χ_L and χ_R known as π -MEMs. On the other hand, with $g < 0.5$, each state has a partner with the same quasienergy, connected by a pair of “0-MEMs”, but these were not a focus of this experiment. These edge modes can be expressed as a sum of Majorana fermions, decaying away from the edge with some localisation length ξ – upon application of a Jordan-Wigner transform, these become non-local operators with long parity strings (see Fig. 2.13). Fig. 2.12(a) shows the time evolution of $\langle \sigma_1^z(t) \rangle$ following a quench from a random σ^z basis product state in the $h_j = 0$ and $g > 0.5$ regime. Looking at the Fourier transform of this quantity, the single-particle modes can be resolved as peaks at the corresponding quasienergies, with two sharp peaks near $\omega = \pi$ due to the π -MEMs. Fig. 2.12(b) then plots this spectrum against g for $L = 6, 12$, and 18 , showing that near $g = 0.5$ there is a large splitting δ between the the two π -MEMs, but this narrows as g gets larger and L increases. On the other hand, there is also a bulk-edge gap Ξ that remains nonzero for all g even with increasing L . If our understanding of SPT edge modes is correct, the splitting Δ should shrink exponentially with L – and this is exactly observed in Fig. 2.12(c), at least for g sufficiently far from the critical point.

Intriguingly, after the integrability-breaking perturbations h_j are turned on (with the h_j distributed i.i.d. in $[-1, 1]$), signatures of the edge modes are preserved in the dynamics: $\langle \sigma_j^z(t) \rangle$ for $j = 1, L$ alternate in sign with slowly-decaying magnitude, even as the same quantity in the bulk rapidly decays to zero. These long lifetimes are attributed to prethermalisation, caused by the bulk-edge gap Ξ depicted in Fig. 2.12(b). Additionally, the quantity $\langle \sigma_1^z(0) C(t) \rangle$ exhibits similar slow exponential decay for certain Pauli strings $C = \left(\prod_{j < k} \sigma_j^x \right) \sigma_k^{y,z}$, which the authors attribute to overlap with the corresponding π -MEM. By calculating the decay rate for a number

of such C the left π -MEM, χ_L , can be reconstructed, and the magnitudes of the coefficients – shown in Fig. 2.13(b) – are found to decay exponentially into the bulk, exactly as predicted by theory. A similar method can also be used to find the right π -MEM, χ_R . The localisation length ξ of this decay was also found by Ref. [187] to diverge approaching the critical point $g = 0.5$, as expected.

2.4.3 Robust topological order at finite temperature

Many quantum technologies require the ability to coherently store quantum information, protected from noise, for long time periods: for example, a quantum computer may need to store intermediate results of calculations. But this is still some way off at present: state-of-the-art quantum memories, even with error-correcting codes, are only able to achieve coherence times of about 2 ms [188]. For this reason, there is considerable interest in mechanisms which can lead to long coherence times, and one obvious candidate is topological order – in particular the edge modes. However, at finite temperatures topological order typically vanishes, and so methods to stabilise it beyond the ground state are needed.

Eq. (2.16) suggests the use of a strong zero mode, but unfortunately, upon adding integrability-breaking interactions (such as a nearest-neighbour $\sum_j \Gamma_2 \sigma_j^x \sigma_{j+1}^x$ or next-nearest-neighbour $\sum_j J_2 \sigma_j^z \sigma_{j+1}^z$ term), the expression becomes vastly more complicated, with a combinatorial increase in terms per order of perturbation theory [189]. Furthermore, the expression no longer converges to an exact strong zero mode, with the error $\|[H, \Psi]\|$ now *growing* after some critical order n^* (which depends upon the model parameters), even in the thermodynamic limit. As a result the perturbative expansion must be truncated at this order, resulting in a strong *almost* zero mode which is only approximately conserved; that is, it has a finite coherence time, with σ_1^z . The expansion is also expected to break down at some order in perturbation theory when J_2/J is rational due to poles in the expansion, which presents further problems: these represent resonant processes enabled by the interaction terms which can decohere the edge mode. One might hope to pick an irrational J_2/J , but since the rationals are dense in the reals, there will always be a pole nearby at a sufficiently high order in perturbation theory. However, with the right choice of J_2/J , this may take a very long time to affect the dynamics. This can also be looked at through the lens of prethermalisation, in which a system relaxes at short times to a state

described by an effective Hamiltonian with potentially very different features to the true Hamiltonian (or Floquet operator), before later approaching thermal equilibrium [190].

Even so, at finite temperatures and with finite interaction strengths, this edge mode is also only able to store *classical* information: while σ_1^z is conserved, to store quantum information phase coherence must be preserved [191]. Another way to put this is that, in terms of logical operators, we need both σ^z and σ^x to have a long-lived qubit. Ref. [28] finds that in a certain spin-1/2 chain with three-site “cluster” operators and $\mathbb{Z}_2 \times \mathbb{Z}_2$ SPT order, a coherent edge qubit is able to survive at infinite temperature when the odd and even sublattices are detuned from each other. Our research in Chapter 3 suggests that quantum scars may be responsible for this behaviour, that these scars display signatures of SPT order, and additionally finds that the scarred manifold exhibits coherence in the bulk of the chain. This points the way towards using quantum scars to preserve topological order at finite temperatures, and potentially using this to construct robust quantum memories. Indeed, projector embedding has been used to insert topologically-ordered scars into otherwise chaotic models [17, 179], but it is not clear if these scars would be stable away from fine-tuned points. The AKLT model, which hosts quantum scars [10] but also topological order [21, 22, 24, 25], would also appear to be a good candidate. However, despite being directly related to the SPT-ordered ground state by a set of raising operators, the scars do not appear to exhibit SPT order. Additionally, the chain hosts a continuous $SU(2)$ symmetry, which is known to prevent MBL [192] – this would suggest that an edge mode could not survive at finite temperatures either.

Clearly, ergodicity breaking through quantum scars can be used to preserve SPT order, which suggests an alternative strategy of using MBL. In MBL systems, excited states resemble the ground states of gapped Hamiltonians, which likewise implies that when the ground state is topologically ordered, the excited states might replicate that order. Indeed, studies have shown that applying quenched disorder to a model which hosts topological order can extend it throughout the spectrum [29–31, 112, 124, 177, 178, 193], and in fact create multiple MBL phases, one for each kind of topological order (including the topologically trivial phase). However, as discussed in Sec. 2.2.4, the stability of MBL in the thermodynamic limit is an

ongoing debate, and in Chapter 4 details our research into the breakdown of MBL near an SPT phase transition, recently reported by a number of small-system exact diagonalisation studies [32–35]. We also show how, in the topologically non-trivial phase, one of the l-bits becomes the SPT strong (almost) zero mode.

Regardless of the mechanism for avoiding thermalisation, it is clear that interactions between ergodicity breaking and quantum order lead to many varied and complex phenomena, with important lessons for condensed matter physics and even potential applications in emerging quantum technologies. It is my hope that this thesis shines some light on this rich area of research.

Chapter 3

Quantum scars in a symmetry protected topological phase

3.1 Introduction

SPT order is usually associated with protected edge modes, degenerate entanglement spectra, and non-local order such as string order or topological entanglement [21–27], but at finite temperature interactions with the bulk cause the edge modes to decohere, and other signatures of SPT order are lost. Despite this, there have been various approaches to preserving topological order at finite energy density. In particular it has been shown that many body localisation (MBL), induced by strong disorder, can protect topological edge modes and other signatures of topological order, up to arbitrarily high energies [29, 30]. However, since a system in an MBL phase cannot sustain bulk transport, the distinguishing features of decoupled edge modes are weakened. More recent work has shown that dimerising a certain spin chain which has SPT order, the so-called ZXZ model, energetically decouples the bulk from the topological edge modes [28], enabling lifetimes that are exponentially long in system size and circumventing this problem.

Following the observation of long-lived oscillations in a Rydberg atom simulator, a new type of ergodicity breaking was discovered: weak ergodicity breaking, caused by a discrete set of highly-atypical eigenstates in the middle of an otherwise chaotic spectrum [9, 18] (see Sec. 2.3). “Towers” of exact scarred eigenstates have also been discovered in the AKLT chain [10, 14], a paradigmatic toy model of symmetry-protected topological (SPT) order [21, 22, 24, 25] (see Sec. 2.4.2), while the Shiraishi-Mori construction [17] allows arbitrary eigenstates to be embedded into chaotic




Figure 3.1: **Top:** AKLT ground state with periodic boundary conditions. Each spin-1 (open ellipses) is split into two virtual spin-1/2 degrees of freedom (filled circles). These spin-1/2s are then paired into singlets (black lines), ensuring that no two adjacent sites have total spin $S = 2$. **Bottom:** AKLT ground state with open boundary conditions. The virtual spin-1/2s at the ends are left unpaired, becoming measurable edge zero modes and leaving the ground state fourfold degenerate.

Hamiltonians at finite energy density, including states with topological order [179, 180, 194]. Given the role that quantum scars play as ETH-violating islands in a chaotic sea of thermal states, usually with highly atypical operator expectation values, it is reasonable to wonder whether they may too be able to preserve signatures of SPT order in systems without fine-tuning at finite temperature, without destroying the distinguishing features of decoupled edge modes. This led to investigations into two different spin chains known to host SPT order: the AKLT model in Sec. 3.2, and the ZXZ or “cluster” model in Sec. 3.3. While we did not find evident of SPT order in excited states in the AKLT model, we found that a dimerised ZXZ model with weak interactions hosts scarred states in a non-integrable phase, with evidence of the preservation of SPT order in excited non-thermal states.

3.2 Lack of SPT order in AKLT scars

The AKLT model, named after Affleck, Kennedy, Lieb, and Tasaki who first formulated it [21, 22], is a spin-1 Heisenberg chain with an additional biquadratic interaction, and has become the paradigmatic example of SPT order in spin chains. The model is specified by,

$$H_{\text{AKLT}} = \sum_j P_{j,j+1}^{(2)}, \quad P_{j,j+1}^{(2)} \propto \vec{S}_j \cdot \vec{S}_{j+1} + \frac{1}{3} \left(\vec{S}_j \cdot \vec{S}_{j+1} \right)^2, \quad (2.13)$$

and has a unique ground state under periodic boundary conditions. The projectors $P_{j,j+1}^{(2)}$ map two spin-1s onto the joint spin-2 subspace. As shown in Fig. 3.1, by splitting each physical spin-1 into two virtual spin-1/2 degrees of freedom and then pairing spin-1/2s from adjacent physical sites into spin singlets, it becomes impossible to form a spin-2 from any two adjacent sites, and so the projectors annihilate this state. Since the projectors are positive semi-definite, this state at $E = 0$ must be the ground state. However, under open boundary conditions, there is a spin-1/2 degree of freedom left unpaired on each end of the system, and left unconstrained by the Hamiltonian. As a result, the ground state becomes four-fold degenerate (since each spin-1/2 has local dimension 2), and the system acquires the edge zero modes associated with SPT order. The ground state also has a two-fold degenerate entanglement spectrum (2.4), with total entanglement entropy a constant $S = \ln 2$, which follows from the fact that any bipartite cut must sever one of these spin singlets. Finally, the ground state may be represented as a translationally invariant MPS with local tensors,

$$A^+ = \sqrt{\frac{2}{3}}\sigma^+, \quad A^0 = -\sqrt{\frac{1}{3}}\sigma^z, \quad A^- = -\sqrt{\frac{2}{3}}\sigma^- . \quad (3.1)$$

Given that $\sigma^\pm(\sigma^z)^n\sigma^\pm = 0$ for any $n \geq 0$, in the S^z basis the ground state is formed of alternating $|+\rangle$ s and $|-\rangle$ s with some number of intervening $|0\rangle$ s – for example, $|+00-0+00-+-\rangle$. This leads to string order in the ground state, which may be quantified by the string order parameter [185],

$$\mathcal{O}_{\text{str}}^\alpha = \lim_{|j-k| \rightarrow \infty} \left\langle S_j^\alpha e^{i\pi \sum_{j \leq \ell < k} S_\ell^\alpha} S_k^\alpha \right\rangle . \quad (3.2)$$

For $\alpha = z$, the exponent counts the balance between $|+\rangle$ s and $|-\rangle$ s, and so it can be seen that any basis state satisfying the above condition will leave a negative contribution, such that $\mathcal{O}_{\text{str}}^z = -1$ in the SPT-ordered phase. Since there is nothing special about the z direction under the Hamiltonian (2.13), this will also hold true for $\alpha = x, y$ or indeed any spin unit vector.

Recently, exact scarred eigenstates were discovered in the AKLT model [10, 14], with the majority of these scars taking the form of a “tower of states” built from the

ground state. These states are given by

$$|\mathcal{S}_{2n}\rangle = \mathcal{N}(Q^\dagger)^n |\text{g.s.}\rangle , \quad (3.3)$$

$$Q^\dagger = \sum_j (-1)^j (S_j^+)^2 , \quad (3.4)$$

where S_j^+ is the spin-1 raising operator on site j , $\mathcal{N}$ is a normalisation constant, and $|\text{g.s.}\rangle$ is an appropriate ground state for the boundary conditions. Each $(S_j^+)^2$ forces the spin at site j to be in the spin-up, $m = 1$ state, and breaks the adjoining singlet bonds. Therefore, each application of $Q^\dagger$ creates a spin-2 magnon excitation with momentum $k = \pi$, and so the state $|\mathcal{S}_{2n}\rangle$ has total spin $S = 2n$. Remarkably, these states also have exact energies $E_{2n} = 2n$ and exhibit log-law entanglement entropies, meaning that they violate the strong ETH. Note that, due to the $\text{SU}(2)$ symmetry of the model, each of these scars in fact forms a multiplet of rotated states, and the construction here gives the highest-weight ($m_z = 2n$) state.

Given the known SPT order in the AKLT ground state, this naturally leads to the question of whether any of these scars also exhibit SPT order. However, we would appear to run into problems even from an analytical perspective. The magnons explicitly break the spin-singlet bonds, which means it is unlikely the scars would have the same two-fold degenerate entanglement spectrum, while also acting on the edge modes. Furthermore, the pattern of alternating $|+\rangle$ s and $|-\rangle$ s is broken, since the magnon operator acts locally to flip spin downs to spin ups. Numerical simulations (not shown here) also failed to show detectable SPT order in the scar states, although I cannot rule out that it persists in some hidden form.

3.3 Quantum Scars in the ZXZ Model

It is a known result that unbroken continuous non-Abelian symmetries – such as the $\text{SU}(2)$ symmetry in the AKLT model – prohibit MBL [151, 152, 192]. The intuitive explanation is that the l-bits transform under representations of the symmetry, promoting the global symmetry to a local one, and thus introducing an exponential degeneracy of eigenstates which is unstable to even an infinitesimal perturbation. In this case, either the symmetry must break or localisation fails, and so localised topological phases protected by non-Abelian symmetries would be impossible. In light of this, having failed to find evidence of SPT order in the excited states of

the AKLT model, it seemed reasonable to expect greater success looking at models whose topological order is only protected by discrete symmetries.

Hence in this section, we study a spin-1/2 chain with three-site “cluster” terms which is known to host a $\mathbb{Z}_2 \times \mathbb{Z}_2$ SPT phase, known as the ZXZ or cluster model [195]. This model is also equivalent in a non-interacting limit to two copies of the transverse field Ising model (TFIM). The ZXZ model initially gained interest in the quantum information literature as a universal resource state for entanglement [196], allowing measurement-based quantum computation protected by symmetry [197, 198], and can be reproduced in optical lattices [28, 199] or as the low-energy subspace of a simple two-body Hamiltonian [200, 201].

However, more recently, the model has provided insight into thermalisation, and in particular how SPT order may be preserved at finite temperature. There have been two successful approaches towards this. By applying quenched disorder and inducing many-body localisation, it was found that the topological edge modes were able to survive up to arbitrarily high energies [30]. However, this technique also arrests bulk transport, masking the distinguishing features of the protected edge mode. More recently, it was shown that dimerising the model, by energetically detuning the cluster terms on odd and even sites respectively, is able to prevent edge mode excitations from being absorbed by the bulk, enabling lifetimes exponentially long in system size [28]. Inspired by this work, we considered whether these long lifetimes might be related to quantum scarring, and whether it would be possible to preserve the bulk given the correct initial conditions.

Fig. 3.2(a) shows a schematic depiction of the model. While it is formulated as a 1D spin chain, it is better viewed in this context as a triangular ladder, with the three-body terms (known as cluster operators) acting on the sites surrounding each plaquette. The legs of the ladder are comprised of the odd and even sublattices respectively. Studying the eigenstates of the ZXZ model in the weakly interacting regime $\lambda = 0.6$, $\Gamma = 0.1$ and $V = 0.05$, we discover a class of non-thermal states in the spectrum, with sub-volume law entanglement entropies and atypical cluster operator expectation values; see Fig. 3.2(b). Notably, these eigenstates appear to come in sets with almost-equal energy spacing, and have one or both sublattices fully polarised. Labelling states by the expectation values of cluster operators, as in




Figure 3.2: (a) The ZXZ Hamiltonian (3.5) describes a spin-1/2 chain dominated by three-site “cluster” terms, $\sigma_{\ell-1}^z \sigma_\ell^x \sigma_{\ell+1}^z$, with a global field $\Gamma \sigma_\ell^x$ at each site. However, it is best thought of as a triangular ladder, with the odd and even sublattices respectively comprising the legs, and an additional integrability breaking term $V \sigma_\ell^x \sigma_{\ell+1}^x$ acting across each rung. (b) With $V \neq 0$, the model is non-integrable, with the spectrum showing level repulsion. Despite this, we find eigenstates throughout the spectrum with sub-volume law entanglement entropies, coexisting with thermal states obeying a volume law. These atypical states have fully-polarised cluster operators on one or both sublattices. (c) It is possible to construct a set of mutual eigenstates of the cluster operators which we label by their eigenvalues on each site, forming a complete basis within each symmetry sector; see Sec. 3.4.1. These are also eigenstates of the $\Gamma = V = 0$ model. (d) Initial states with fully-polarised sublattices, like the atypical states in (b), show unusual dynamics including preservation of the Loschmidt echo $|\langle\psi(t)|\psi\rangle|^2$, while states without this feature rapidly thermalise with cluster operators approaching the thermal ensemble average. Here, we choose $L = 12$, $\Gamma = 0.1$, $V = 0.05$, $\lambda = 0.6$.

Fig. 3.2(c), we then show that states with fully-polarised sublattices exhibit atypical dynamics, while states without this feature decayed to an equilibrium value. This is demonstrated in Fig. 3.2(d), showing the Loschmidt echo $|\langle\psi(t)|\psi\rangle|^2$ tending quickly to zero for a typical state but to a large constant value for the state $|1\bar{1}\bar{1}\bar{1}\bar{1}\bar{1}\bar{1}\bar{1}\bar{1}\bar{1}\bar{1}\bar{1}\bar{1}\bar{1}\rangle$.

The rest of this section is structured as follows. First, in Sec. 3.3.1, we formally introduce the ZXZ model and its known properties, and then Sec. 3.3.2 contains a brief account of the techniques used to study the model. Having set up the problem and methods, in Sec. 3.3.3 we present the major results, namely the presence of a

tower of quantum scars in the ZXZ model with non-thermal properties and signatures of SPT order. Sections 3.3.4 and 3.3.5 cover two important areas which are necessary for a complete understanding of the model: the approach to full quantum chaos and Wigner-Dyson statistics in the former, and the unusual dynamics of the model in the latter. Finally, we round off the section with a discussion of the results, their implications and potential directions of future research.

3.3.1 Model

The ZXZ Hamiltonian for a chain with L lattice sites and open boundary conditions (OBC) is given by

$$\mathcal{H} = \sum_{\ell=1}^{L/2-1} (\lambda K_{2\ell} + K_{2\ell+1}) + \Gamma \sum_{\ell=1}^L \sigma_{\ell}^x + V \sum_{\ell=1}^{L-1} \sigma_{\ell}^x \sigma_{\ell+1}^x, \quad (3.5)$$

where the cluster operators K_{ℓ} are given by,

$$K_{\ell} = \sigma_{\ell-1}^z \sigma_{\ell}^x \sigma_{\ell+1}^z, \quad (3.6)$$

and $\sigma_{\ell}^x, \sigma_{\ell}^z$ are the Pauli spin operators at site ℓ . While this nominally describes a 1D spin chain, the model can be best thought of as a triangular ladder with the odd and even sublattices each corresponding to a leg, as was shown earlier in Fig. 3.2(a). The cluster operators link three sites around each triangular plaquette, while the integrability-breaking interaction terms $\sigma_{\ell}^x \sigma_{\ell+1}^x$ form the rungs of the ladder. A local field σ_{ℓ}^x is also applied to each site. With just these three terms, $\mathcal{H}$ has a $\mathbb{Z}_2 \times \mathbb{Z}_2$ symmetry, given by spin inversion on odd and even sites respectively, and additionally when $\lambda = 1$ there is an additional $\mathbb{Z}_2$ parity symmetry corresponding to the transformation $\ell \rightarrow L - \ell + 1$ (which swaps odd and even sites). Note that, throughout this work, it is understood that while $\mathcal{H}$ has L lattice sites with subscripts $1 \leq \ell \leq L$, due to OBC there are only $L-2$ cluster terms K_{ℓ} labelled by $2 \leq \ell \leq L-1$.

The cluster operators K_{ℓ} are mutually commuting, and hence there exists a basis of states which are mutual eigenstates of all the K_{ℓ} , labelled by their eigenvalues on each site, which we call the cluster basis (see Sec. 3.4.1). These states are exact eigenstates of the model when $\Gamma = V = 0$, and when $\lambda > 0$ this means that the

unique ground state in each symmetry sector is the state $|\text{g.s.}\rangle$ such that,

$$\langle \text{g.s.} | K_\ell | \text{g.s.} \rangle = -1, \quad 2 \leq \ell \leq L-1. \quad (3.7)$$

This also implies that the spectrum of $\mathcal{H}$ is trivial when $\Gamma = V = 0$, and that the model is integrable. Surprisingly, the model also remains integrable for $\Gamma \neq 0$, and it is only the two-body interaction term which breaks integrability. It can be shown that the cluster operators are equivalent through a Jordan-Wigner transform to a fermionic next-nearest neighbour interaction [28, 29, 195, 202] such that when $V = 0$, the odd and even sublattices become entirely disconnected. This means that the even sites and odd sites become entirely disconnected, and the model can be separated into two copies of the transverse-field Ising model (TFIM), $\mathcal{H}(V = 0) \sim \mathcal{H}_{\text{TFIM}}^e(\lambda_e = \lambda, \Gamma) + \mathcal{H}_{\text{TFIM}}^o(\lambda_o = 1, \Gamma)$, which is well known to be integrable [203],

$$\mathcal{H}_{\text{TFIM}}^{e/o}(\lambda, \Gamma) = \lambda \sum_{\ell \in e/o} \sigma_\ell^z \sigma_{\ell+2}^z + \Gamma \sum_{\ell \in e/o} \sigma_\ell^x, \quad (3.8)$$

where the sums in Eq. (3.8) should be understood as either running over the even (e) or odd (o) lattice sites. However, turning on nearest-neighbour interactions with nonzero V breaks integrability, and couples the two chains.

The ZXZ model (3.5) exhibits symmetry-protected topological order, protected by the $\mathbb{Z}_2 \times \mathbb{Z}_2$ spin-flip symmetry [195]. As a result, the model hosts robust boundary degrees of freedom at zero temperature, with other signatures of this phase including a four-fold degenerate entanglement spectrum and string order in the ground state. Such edge modes typically decohere at finite temperature, due to interaction with thermal excitations in the bulk, but Ref. [30] demonstrated that they can be stabilised by inducing MBL through strong quenched disorder. In this case, the life time of the zero modes grows exponentially with the size of the system, such that the boundary qubit remains coherent on indefinite time scales in the thermodynamic limit.

Additionally, it has recently been shown in Ref. [28] that if $\lambda \neq 1$ in Eq. (3.5), such pre-thermal boundary modes can also survive at infinite temperature without disorder. (Note that while $\lambda \neq 1$ preserves the SPT-protecting $\mathbb{Z}_2 \times \mathbb{Z}_2$ spin-flip symmetry, it breaks the $\mathbb{Z}_2$ “swap” symmetry between the legs.) This choice leads to a dimerisation of the chain such that each leg of the ladder has a different cost

for excitations, resulting in an energetic detuning. Crucially, this detuning ensures that excitations on one leg cannot be resonantly converted into excitations on the other leg, which prevents the bulk from absorbing the boundary modes. Just like in Ref. [30], the lifetime of the pre-thermal edge modes is exponential in system size and both the x & z components are preserved, allowing the edge mode to store quantum information.

To elaborate, this exponentially long-lived qubit is explained by two conjugate almost strong zero modes (see Sec. 2.4.2) localised to the edge of the system. Under the mapping to two copies of the TFIM, in the $V \neq 0$ regime, the system is now dual to one described by the Hamiltonian,

$$\mathcal{H}_{\text{TFIM}}^e(\lambda_e \lambda, \Gamma) + \mathcal{H}_{\text{TFIM}}^o(\lambda_o = 1, \Gamma) + V \sum_{\ell=1}^{L-1} \sigma_{\ell}^x \sigma_{\ell+1}^x. \quad (3.9)$$

In the ferromagnetic phase [$\lambda > \Gamma$ in Eq.(3.8)], the non-interacting TFIM is known to contain an exact strong zero mode or SZM (2.16) at the edge [189]. With the addition of interaction terms this no longer converges in the thermodynamic limit, becoming instead an almost strong zero mode which must be truncated at some finite order of perturbation theory, and furthermore when parameters take rational values with small denominators, the lifetime becomes extremely small. Additionally, the SZM in the interacting TFIM only conserves the z component of the edge spin, making it a store of classical and not quantum information. However, the ZXZ model contains *two* copies of the TFIM (3.9). As explained by Ref. [28], the respective edge modes are σ_1^z and σ_2^z to leading order, but these are not conjugate to each other and so do not form a qubit. However, the global spin-flip symmetry operator on odd sites, $G_o = \prod_{\ell \in \text{odd}} \sigma_{\ell}^x$, commutes with the Hamiltonian, and so we can multiply the second edge mode by this to obtain $\sigma_2^z G_o$ which *is* conjugate to σ^x . Upon mapping back to the original ZXZ model, these become $\Sigma_z = \sigma_1^z$ and $\Sigma_x = \sigma_1^x \sigma_2^z$ respectively, and it is these that form the leading-order components of the conserved qubit.

In the Ising picture, when $V = 0$, the Ising chains are decoupled, and neither SZM can decay without absorbing or emitting a domain wall. These domain walls incur an energy cost of $2\lambda_{e/o}$, making this an off-resonant process. However, with nonzero V , these domain walls can hop between the chains, which means that when

$\lambda = 1$ the two edge modes can resonantly exchange excitations – causing both to rapidly decohere. However, if $\lambda \neq 1$, then there is an additional energy cost to convert a domain wall from one chain to another, obstructing this process – but with rational λ , higher order processes will still lead to resonances. This is made manifest when perturbatively expanding the two SZMs starting from Σ_x and Σ_z , order by order – here, the expansions acquire poles corresponding to rational values of λ . Since the rationals are dense in the real numbers, even an irrational choice of λ will lead to large coefficients in the SZM expansions at some order, limiting the maximum coherence time of the edge qubits – Ref. [28] found an optimal value of $\lambda = 0.6$ for system sizes accessible by exact diagonalisation.

In this work, we looked at the ZXZ model in the weakly interacting regime $\Gamma = 0.1$, $V = 0.05$, similar to the values chosen in Refs. [28, 30], and likewise set $\lambda = 0.6$. $\mathcal{H}$ becomes non-integrable with this parameter choice, however for the system sizes studied here, the model is not strongly thermalising, sitting in an intermediate regime between full quantum chaos and integrability. This can be seen in the comparatively broad distributions of entanglement entropies (Sec. 3.3.3.3) and cluster operator expectation values (Sec. 3.3.3.4), as well as in level spacing statistics (Sec. 3.3.4).

3.3.2 Methods

Results were obtained using exact diagonalisation (ED) methods, in which the Hamiltonian on a finitely-sized system is written down exactly as a matrix in some computational basis, and then diagonalised to find the full spectrum of energy eigenstates and eigenvalues. In principle, this allows the calculation of arbitrary observables, at any energy. However, the technique comes at the cost of the computational resources required (in particular memory) scaling polynomially with Hilbert space dimension. Since this in turn grows exponentially with system size, this strongly limits the maximum system size that can be studied even in principle, and strategies are required to either split up the Hilbert space into smaller disconnected sectors (such as by taking advantage of symmetries in the problem [204, 205]), or to obtain only a limited selection of eigenstates in return for reduced memory usage. In this case, we were able to use the $\mathbb{Z}_2 \times \mathbb{Z}_2$ symmetry to reduce the Hilbert space dimension by a factor of 4, in turn enabling calculations of system sizes up to $L = 18$. The

specific software I used was QuSpin 0.3.4 [206], an open-source Python package for fast, multi-threaded ED and quantum dynamics calculations.

The general procedure for this project was as follows. We started our investigations by exploring the transition between integrable and chaotic behaviour in the model (see Sec. 3.3.4), in order to determine appropriate parameter choices such to ensure the model was in a non-integrable regime. Following this, I diagonalised the Hamiltonian for these parameter choices, and investigated the distributions of various operator expectation values and the bipartite entanglement entropy among the energy eigenstates: the aim here was to discover states with atypical values which would be candidates for quantum scars. This led to the discovery of the towers of states detailed in this chapter, which were then carefully characterised using further ED techniques to obtain the data presented below. Jonas Richter additionally contributed simulations of the dynamics, in order to probe the behaviour of the model following quenches from the scarred manifold.

3.3.3 Eigenstate Analysis

Quantum many-body scars are eigenstates of an otherwise ergodic Hamiltonian which violate the ETH, meaning that they are non-thermal, despite having finite energy density. This should be contrasted with integrable and many-body localised systems, in which every eigenstate violates the ETH, and correspondingly with fully chaotic systems, in which every state with non-zero energy density is thermal. In this subsection, we will first relate excitations in the TFIM to those in the ZXZ model, showing that they are equivalent in the regime of perturbatively small Γ and zero V . We will then present evidence that a set of non-thermal eigenstates exists in the ZXZ model, with some occupying the middle of the spectrum, and that these are well-approximated by a generalisation of these excitations. These states form a “tower” of states connected by a raising operator (c.f. Refs. [10–16]), and are representable as matrix product states (MPS). Finally, we will characterise these eigenstates in detail, showing that they have bipartite entanglement entropies obeying a sub-volume law and exhibit atypical expectation values of local cluster operators. Additionally, we will show that these states may preserve a four-fold degenerate entanglement spectrum, a key signature of $\mathbb{Z}_2 \times \mathbb{Z}_2$ SPT order.

3.3.3.1 Low-lying excitations

The non-interacting ($V = 0$) Hamiltonian is separable into two copies of the transverse-field Ising model, corresponding to the odd and even sublattices. As such, the low-energy excitations of the model should be equivalent to the excitations of the TFIM. In the regime we investigate, the ZXZ model maps onto the ferromagnetic phase of the TFIM – that is, where $\lambda \gg \Gamma$ in Eq. (3.8). In this phase of the TFIM, excitations are given by domain walls between two regions of aligned spins, while the σ_j^x term causes these walls to hop, leading to delocalised excitations [29]. The domain walls in the TFIM map onto cluster excitations in the ZXZ model [28, 29, 195], and so we can apply the same techniques as used to solve the TFIM. Consider the cluster raising (lowering) operator, which creates (destroys) an excitation at a given site,

$$K_\ell^\pm = \frac{1}{2} (\sigma_\ell^z \mp i \sigma_{\ell-1}^z \sigma_\ell^y \sigma_{\ell+1}^z) . \quad (3.10)$$

By expressing a bulk σ_j^x operator in terms of these operators,

$$(K_{j-1}^+ + K_{j-1}^-) K_j (K_{j+1}^+ + K_{j+1}^-) = \sigma_j^x , \quad (3.11)$$

it is clear that the σ_j^x term causes two next-nearest neighbour cluster operators to flip. This will cause an isolated cluster excitation to hop two sites at a time, at zero energy cost, or create and destroy a pair of excitations while changing the energy of the state. Considering a perturbatively small Γ , the zero-energy parts of this term mix the cluster states under degenerate perturbation theory, with a first order effect in the energy, while the other parts act only to higher orders. Therefore, when this condition is satisfied, the low-lying eigenstates in the ZXZ model with periodic boundary conditions should resemble delocalised cluster excitations.

If we are in a regime such that $\Gamma \ll \lambda \leq 1$, then we can work in the space of a fixed number of cluster excitations, which greatly simplifies calculations. Assume that there is just one excitation and that it is on the odd sublattice. The effective Hamiltonian in this space is then,

$$\mathcal{H}_{\text{eff}} |2\ell + 1\rangle = |2\ell + 1\rangle - \Gamma (|2\ell - 1\rangle + |2\ell + 3\rangle) , \quad (3.12)$$

where $|\ell\rangle = K_\ell^+ |\text{g.s.}\rangle$ is a single excitation localised at site ℓ . This can be solved by introducing delocalised momentum states $|k\rangle$ via a Fourier transform,

$$|k\rangle = \sqrt{\frac{2}{L}} \sum_{\ell=0}^{L/2-1} e^{+ik\ell} |2\ell+1\rangle, \quad (3.13)$$

$$|2\ell+1\rangle = \sqrt{\frac{2}{L}} \sum_{k=0}^{L/2-1} e^{-ik\ell} |k\rangle, \quad (3.14)$$

where k can take the values $2\pi\eta/(L/2)$ for integer η , $0 \leq \eta < L/2$. Substituting into Eq. (3.12), and noting that the states $|k\rangle$ are mutually orthogonal, we obtain:

$$\mathcal{H}_{\text{eff}} |k\rangle = (1 - 2\Gamma \cos k) |k\rangle. \quad (3.15)$$

This is the same as for an excitation in the transverse-field Ising model, showing again the mapping between the ZXZ model and two copies of the TFIM. Note that with minor changes, similar excitations can be shown to exist as holes in a fully-excited sublattice, and also when the other sublattice is fully excited.

This result is only strictly valid with periodic boundary conditions, perturbatively small Γ , and $V = 0$. Yet with some modifications, we find that these states also convincingly approximate low-lying excitations in the ZXZ model with open boundary conditions and non-zero V . More surprisingly, we also find that states with multiple such excitations are still close to a generalisation of the above result.

3.3.3.2 Approximate tower of states

Working with OBC, the sum over ℓ in Eq. (3.13) now starts at $\ell = 1$, as the first odd cluster site is $K_{2\ell+1} = K_3$, c.f. Eq. (3.5). Moreover, the allowed values of k are $k = 2\pi\eta/(L/2 - 1)$, $0 \leq \eta < L/2 - 1$, and the normalisation changes. With these changes, we find that the states $(|+k\rangle \pm |-k\rangle)/\sqrt{2}$ each have a high overlap with a particular eigenstate of the model, for every allowed value of k , even with $V > 0$.

To generalise this result to multiple excitations, we will introduce an operator which creates a single delocalised excitation at a time. To start with, consider the cluster raising (lowering) operator (3.10), which creates (destroys) an excitation at a given site. Using this and focusing on excitations at zero momentum ($k = 0$), we

then construct an operator acting upon the entire chain by,

$$\mathcal{O} = \sum_{\ell=1}^{L/2-1} K_{2\ell+1}^+ . \quad (3.16)$$

We then apply the operator repeatedly to the ground state $|\text{gs}\rangle$ (3.7) of the clean ZZZ model to produce a tower of states [10–16]. That is, we generate a set of states $|T_j\rangle$ given by,

$$|T_j\rangle = \frac{\mathcal{O}^j |\text{gs}\rangle}{|\mathcal{O}^j |\text{gs}\rangle|} . \quad (3.17)$$

Because the operator (3.16) produces a single delocalised excitation, each set of states has an energy spacing of approximately 2 – this spacing is exact in the $\Gamma = V = 0$ model. In the insets of Fig. 3.3(a) and (b), we show that the states $|T_j\rangle$ have large overlaps with certain eigenstates of $\mathcal{H}$ in both the non-interacting ($V = 0$) and the interacting ($V = 0.05$) case.

In the following, we label those eigenstates $|n\rangle$ which maximise the overlap $|\langle n|T_j\rangle|^2$ for a given tower index j by $|S_j\rangle$. In particular, the comparatively simple structure of the $|T_j\rangle$ already indicates that the eigenstates $|S_j\rangle$ might exhibit atypical properties, which we will analyse in more detail below. Note, however, that $|T_7\rangle$ in fact has large and comparable overlaps with a pair of adjacent eigenstates, but it is only the one with the larger overlap that we label $|S_7\rangle$. This feature of $|T_7\rangle$ may indicate that the corresponding scar state loses stability for increasing L , as this was not observed for any of the $|T_j\rangle$ in the $L = 14$ case and is only weakly apparent for $L = 16$ for certain j .

If an initial state $|\psi\rangle$ can be found whose spectral decomposition is dominated by the scarred eigenstates $|S_j\rangle$, such a $|\psi\rangle$ should yield periodic oscillations in time. As shown in Fig. 3.3(a) and (b), such a state indeed exists and can be constructed by applying cluster-lowering operators on the even cluster sites to a spin-basis product state (for details, see Appendix 3.4.1). The resulting state $|\psi_{1/0}\rangle$ is a superposition of cluster basis states and has expectation values,

$$\langle \psi_{1/0} | K_{2\ell} | \psi_{1/0} \rangle = -1, \quad \langle \psi_{1/0} | K_{2\ell+1} | \psi_{1/0} \rangle = 0 , \quad (3.18)$$

although it is important to note that these expectation values alone do not define




Figure 3.3: [(a), (b)] Squared overlaps of the eigenstates of $\mathcal{H}$ with the state $|\psi_{\bar{1}/0}\rangle$, plotted against energy, for $L = 18$, in the subspace with positive spin-flip symmetry on both sublattices. Panel (a) shows the results for the non-interacting case ($V = 0$) and (b) for the interacting case ($V = 0.05$). The highlighted eigenstates $|S_j\rangle$ exhibit a significant and atypically high overlap with $|\psi_{\bar{1}/0}\rangle$, and can be well approximated by the tower of states $|T_j\rangle$ [see insets to panels (a) and (b)]. (c) Fourier-transformed Loschmidt echo of the state $|\psi_{\bar{1}/0}\rangle$ for system sizes from $L = 12$ up to $L = 20$, with $V = 0.05$. The inset shows the original time-domain data. [Other parameters: $\lambda = 0.6$, $\Gamma = 0.1$]

this state uniquely. In Fig. 3.3 (c), we show the Fourier-transformed Loschmidt echo $\mathcal{L}(\omega) = \mathcal{F}[\mathcal{L}(t)]$ of the state $|\psi_{\bar{1}/0}\rangle$. As expected from the high overlap with the almost equidistant states $|S_j\rangle$, $\mathcal{L}(\omega)$ exhibits peaks at frequencies ω which are multiples of 2. For increasing system size L , the spectral contributions of peaks at higher ω become slightly less pronounced. Correspondingly, we find that the revivals of $\mathcal{L}(t)$ in time [see inset of Fig. 3.3 (c)] become less distinct for increasing L .

Note that in addition to the $|T_j\rangle$ discussed above, there are in fact three other towers starting from the ground state of each symmetry sector. These can be

obtained by replacing the raising operators in Eq. (3.16) with lowering operators and starting from the most excited state, or by applying the operators to even sites, or by a combination of both. Physically, these correspond to the transformations $K_\ell \rightarrow -K_\ell$ and $K_{2\ell} \leftrightarrow K_{2\ell+1}$ respectively. Moreover, while the operator $\mathcal{O}$ produces a zero-momentum excitation, we find that generalisations to $k \neq 0$ as well as to half-integer values of η seem to also yield good approximations to eigenstates of $\mathcal{H}$ (see Appendix 3.4.3 for details). Intriguingly, if we start the construction of the tower from some state besides the ground state (e.g. adding odd excitations to an initial state with a single excitation on the even sites), the above results do not hold, i.e., the states generated are poor matches to eigenstates of $\mathcal{H}$.

Eventually, let us mention one technical detail. Specifically, the operator $\mathcal{O}$ in Eq. (3.16) changes the $\mathbb{Z}_2$ spin-flip symmetry on the odd sublattice such that the tower produced starting from a state in one symmetry sector will alternate between that and one other symmetry sector. For computational convenience, we choose to shift the entire tower into the same symmetry sector. This is achieved by applying the operator σ_1^z after every application of the operator $\mathcal{O}$, as σ_1^z anticommutes with the spin flip symmetry on odd sites, but commutes with all cluster operators. As the states $|T_j\rangle$ in each symmetry sector are equivalent, this should not affect the results, and we have checked this by performing the same analysis in different symmetry sectors.

3.3.3.3 Entanglement entropy

Given a state $|\psi\rangle$, its entanglement entropy for a bipartition into subsystems A and B is given by,

$$S_A = -\text{Tr}[\rho_A \ln \rho_A] , \quad \rho_A = \text{Tr}_B\{|\psi\rangle\langle\psi|\} , \quad (2.4)$$

where ρ_A is the reduced density matrix for a subsystem A . For an eigenstate of a Hamiltonian obeying the ETH, one generally expects that S_A scales with the system size, and at infinite temperature it approaches the Page value [207], which is the average entropy for a random pure state. A state with an extensive entanglement entropy is said to be obeying a volume-law, while the ground state of gapped systems is always area-law [208]. Surprisingly, states with sub-extensive entanglement entropies have been found even at finite energy densities in a number of otherwise chaotic models, now usually referred to as quantum many-body scars (see Sec. 2.3).




Figure 3.4: (a) Half-system bipartite entanglement entropy of the eigenstates of the non-interacting ($V = 0$) Hamiltonian, plotted against energy, for $L = 18$ in the subspace with positive spin-flip symmetry on both sublattices. Darker colours indicate a greater number of states at that value. Additionally, the tower of states $|T_j\rangle$ and their corresponding eigenstates $|S_j\rangle$ are highlighted by crosses. (b) The same as (a), but for the interacting ($V = 0.05$) model. (c) The scaling of entanglement entropy with subsystem size in the non-interacting model is shown for the states $|S_j\rangle$. A typical thermal eigenstate (dotted black line) exhibiting a volume law is shown for comparison. (d) The same as (c), but for the interacting ($V = 0.05$) model. [Other parameters: $\lambda = 0.6$, $\Gamma = 0.1$]

Figures 3.4(a) and (b) show the half-system bipartite entanglement entropy for the eigenstates of the ZXZ model (3.5) when $L_A = L/2$, in both the non-interacting ($V = 0$) and the interacting ($V = 0.05$) models. While the entanglement entropies take a broad range of values in both cases, with $V = 0.05$ and especially towards the middle of the spectrum, the distribution becomes strongly skewed towards higher entropies, and most of the states become volume law.

However, some states retain much lower entropies, and this is especially true for the eigenstates $|S_j\rangle$ which we identified in Sec. 3.3.3.2 as exhibiting a large overlap with the tower of states $|T_j\rangle$. In particular, the $|S_j\rangle$ have very low entanglement entropies in the non-interacting case, and most retain these when interactions are turned on. Surprisingly, some of the $|S_j\rangle$ have even lower entropies than their approximations $|T_j\rangle$. On the other hand, some of the states $|S_j\rangle$ with larger j do

attain significantly higher entanglement entropies in the interacting case, though still small compared to most states nearby in energy: we attribute this to the large number of excitations in these states, which increases their complexity and provides more ways for the interaction V to destabilise them. In addition to the half-chain entanglement entropy, Fig. 3.4(c) and (d) show S_A versus subsystem size L_A for the states $|S_j\rangle$. The data show that the entanglement of the $|S_j\rangle$ scales sub-extensively both for $V = 0$ and $V = 0.05$, except for $|S_7\rangle$ where this is not clear. This is in contrast to eigenstates which are nearby in energy, which obey a clear volume law.

In fact, we show in Appendix 3.4.2 that the exact states $|T_j\rangle$ may be represented by a matrix-product state (MPS) of bond dimension $\chi = 4[\min(j+1, L/2-j)] \leq L+4$. Since the entanglement entropy of an MPS is at most $\ln \chi$, this places a logarithmic bound on their entropy. We note, however, that it does not necessarily follow that the eigenstates $|S_j\rangle$ will obey this bound as well.

3.3.3.4 Distribution of cluster excitations

The ETH predicts that the expectation values of (local) physical operators, evaluated with respect to individual eigenstates of chaotic Hamiltonians $\mathcal{H}$, should form a smooth function of energy and agree with the microcanonical ensemble average for that operator. As a result, a distinguishing feature of ETH-violating eigenstates is a significant departure of these expectation values from the energy-resolved average. In the case of the ZZZ model, it is instructive to consider the expectation values of the cluster operators K_ℓ for the eigenstates of the model, taking particular note of the values for the eigenstates $|S_j\rangle$.

Figures 3.5(a) and (b) show the histogram of matrix elements $\langle K_\ell \rangle = \langle n|K_\ell|n \rangle$ for two central sites in the chain with $L = 16$, for the non-interacting and interacting Hamiltonians respectively. The states $|S_j\rangle$ are highlighted. They show that in the non-interacting case, the expectation values for the central two sites fall close to a discrete set of values with little systematic dependence on energy. This discretisation might be explained by proximity to the $\Gamma = V = 0$ point at which the K_ℓ are constants of motion. Once interactions are turned on, the distribution of $\langle K_\ell \rangle$ becomes smoother, though there is still substantial variation. However, some states stay close to their original values, and this is particularly prominent for the states $|S_j\rangle$ which clearly deviate from the microcanonical average.




Figure 3.5: (a) $\langle K_\ell \rangle$ at sites $\ell = 8, 9$ for energy eigenstates of the non-interacting ($V = 0$) model, $L = 16$, in the subspace with positive spin-flip symmetry on both sublattices. Darker colours indicate a greater number of states at that value. Circular markers indicate the states $|S_j\rangle$. (b) The same, but for the interacting model ($V = 0.05$). (c) Probability distribution of $\langle K_{\ell_o} \rangle$ for eigenstates in a narrow window around the state $|S_{L/2-2}\rangle$, relative to the microcanonical ensemble average $\overline{K_{\ell_o}}$. $\ell_o = 7, 9, 9$ for $L = 14, 16, 18$ respectively. The values for the states $|S_{L/2-2}\rangle$ are indicated by open circles with vertical dashed lines, and we include data from all four symmetry sectors. (d) $\langle K_\ell \rangle$ v.s. ℓ , for the state $|S_{L/2-2}\rangle$ and a volume-law state adjacent to it in energy, for $L = 16$. [Other parameters: $\lambda = 0.6$, $\Gamma = 0.1$]

For a more refined analysis, Fig. 3.5(c) looks at the distribution of expectation values on a central odd site ℓ_o , focusing on a narrow energy window centred around the penultimate tower state $|S_{L/2-2}\rangle$. This state was chosen as it is closest to zero energy for the chosen parameters. While the distribution appears to narrow down with increasing system size L , we find that the expectation values $\langle S_j | K_{\ell_o} | S_j \rangle$ remain distinct outliers for all L .

The behaviour of the cluster excitations is explored further in Fig. 3.5(d) by looking at their expectation values for particular states and their variation in space. On the one hand, the expectation value $\langle n | K_\ell | n \rangle$ for an eigenstate $|n\rangle$ directly adjacent to $|S_j\rangle$ is fairly uniform and thermal (as it is the case for the majority of states once interactions are turned on). However since eigenstates $|S_j\rangle$ are well approximated by the tower states $|T_j\rangle$, they have similar cluster expectation values. In particular, this means that the values between sublattices differ substantially with the even sublattice being almost fully polarised.

3.3.3.5 Signatures of SPT order

SPT order is usually associated with properties of the ground state of a system, and in normal circumstances lost at any finite energy density. Despite this, it has been shown that eigenstates in a many-body localised system can remain in a sharply defined topological phase, even at infinite temperatures [29, 30, 177]. More recently, it has been shown that quantum many-body scars embedded into the spectrum of a topologically ordered model can themselves have topological order, despite having a finite energy density relative to the ground state [179, 180, 194]. It is therefore natural to wonder whether the nonthermal states $|S_j\rangle$ discussed in this work might also retain signatures of topological order. To this end, we study the so-called entanglement spectrum [23, 24], defined here as the eigenvalues φ_α of an “entanglement Hamiltonian”,

$$H_{\text{ent}} = -\ln(\rho_A) , \quad (3.19)$$

and we choose subsystem A to be the first $L/2$ sites of the system. (Note that an alternate convention is to use the eigenvalues of ρ_A directly.)

A consequence of SPT order in the ZXZ model (3.5) is that the entanglement spectrum is four-fold degenerate in the ground state, as long as Γ and V are chosen




Figure 3.6: Ordered entanglement energy spectra φ_α [c.f. Eq. (3.19)] of selected exemplary states in $\{|S_j\rangle\}$ (crosses). For comparison, we also show the entanglement spectrum of a typical eigenstate close to zero energy (open circles). We have $L = 18$, $\Gamma = 0.1$, $V = 0.05$ in all cases.

such that $\mathcal{H}$ remains in the $\mathbb{Z}_2 \times \mathbb{Z}_2$ SPT phase. This statement is also true of *every* eigenstate in the $\Gamma = V = 0$ model. However for nonzero Γ , V , the entanglement spectrum of a state in the middle of the energy spectrum should not have this degeneracy.

In spite of this, we find signs of this degeneracy for all of the states in the tower, $|S_j\rangle$. Figure 3.6 gives illustrative examples, as follows. As $|S_0\rangle$ is of course also the ground state, the four-fold degeneracy is clear throughout the entanglement spectrum, and there is a large gap between the “ground” quadruplet and the next set of values. This serves as a benchmark for the behaviour of the other $|S_j\rangle$. States $|S_1\rangle$ and $|S_2\rangle$ still retain a four-fold degeneracy, though with smaller gaps, and a slight breakdown of this degeneracy at higher entanglement energies. Moreover, even though $|S_7\rangle$ is close to the middle of the spectrum, the first dozen φ_α clearly form well-separated quadruplets. Finally even the highest state in the tower, $|S_8\rangle$ keeps the degeneracy in the first quadruplet with a gap almost comparable to that of the ground state $|S_0\rangle$. This may be a signature that these nonthermal states retain SPT order. In particular, the entanglement spectrum of the $|S_j\rangle$ is in stark contrast to a typical state in the spectrum (in this case, chosen to have close to zero energy), which shows no signs of the four-fold degeneracy.

We leave it to future work to study this finding in more detail, e.g., by looking at other indicators such as the topological entanglement entropy or the appropriate string order parameter [209].




Figure 3.7: Distribution of $\tilde{r}$ for $L = 18$, $\Gamma = 0.1$, $V = 0.05$, together with expected distributions for different level statistics. Level spacings are calculated per symmetry sector, in the middle 1/3 of energies, but values from all sectors are included. **(b)** Finite size scaling of the mean value $\langle \tilde{r} \rangle$, for two different choices of Γ and V . The dashed horizontal lines indicate the values corresponding to Poisson, semi-Poisson, and GOE level statistics.

3.3.4 Level statistics and onset of chaos

A useful indicator of whether a system is integrable or chaotic is given by the ratio of adjacent level spacings $\tilde{r}$ [67, 68],

$$\tilde{r} = \frac{\min\{\Delta_n, \Delta_{n+1}\}}{\max\{\Delta_n, \Delta_{n+1}\}}, \quad (3.20)$$

where $\Delta_n = E_{n+1} - E_n$ is the spacing between consecutive energy levels. An integrable system, with extensively many conserved quantities, reduces to a system of non-interacting particles. As a result, its many-body energy levels are independent of one another, such that when considering a macroscopic number of degrees of freedom it is as if the levels have been drawn from a random distribution. This means that the level *spacings* follow a Poisson distribution [210], with no level repulsion observed. On the other hand, for a non-integrable “chaotic” system, it has been shown that the levels will remarkably follow a similar distribution to that of a random matrix drawn from a Gaussian ensemble [65, 66], e.g. the Gaussian orthogonal ensemble (GOE) when $\mathcal{H}$ has real entries, even though the Hamiltonian of the system is typically very non-random. In this case the energies show clear level repulsion. Such “chaotic” systems are expected to obey the ETH [5, 211].

In Fig. 3.7(a), the probability distribution $P(\tilde{r})$ of the ZXZ model is shown for the weakly interacting regime $\Gamma = 0.1$, $V = 0.05$. Data is obtained from the central

third of the spectrum of $\mathcal{H}$ with system size $L = 16$. We find that $P(\tilde{r})$ exhibits clear level repulsion and is inconsistent with a Poissonian distribution. Thus, the small but nonzero V is sufficient to break the integrability of $\mathcal{H}$. However, comparing $P(\tilde{r})$ to the theoretically expected GOE distribution, we also observe distinct deviations. Instead, it appears that $P(\tilde{r})$ is much better described by a so-called semi-Poisson distribution [212], indicating that while $\mathcal{H}$ is non-integrable, full quantum chaotic behaviour is absent. This is consistent with our findings from Figures 3.4 and 3.5, which is that there are many states throughout the spectrum of $\mathcal{H}$ which appear to be not entirely in accord with the predictions of the ETH.

In order to analyse the onset of quantum chaos further, Fig. 3.7(b) shows the mean value $\langle \tilde{r} \rangle$ versus system size L . In agreement with our previous observation in Fig. 3.7(a), we find that $\langle \tilde{r} \rangle$ does not reach the expected value $\langle \tilde{r} \rangle_{\text{GOE}} \approx 0.53$, although we can not rule out that this value is eventually reached for even larger values of L . In contrast, if we consider slightly larger values of the perturbations, $V = 0.1$ and $\Gamma = 0.2$, we find a convincing agreement $\langle \tilde{r} \rangle \approx \langle \tilde{r} \rangle_{\text{GOE}}$ for the largest system with $L = 16$. This suggests that full quantum chaos is restored in this parameter regime. Importantly, however, we note in Sec. 3.3.5 that the atypical dynamics of some initial states can still be qualitatively observed for these larger values of Γ and V .

3.3.5 Non-equilibrium dynamics

Note that, while they were obtained in collaboration and build off of my own work on quantum scars, the calculations in this section were performed by Jonas Richter. However, I include them in order to provide a more complete picture of scarring in the ZXX model.

A key signature of quantum scarring is the strong dependence of dynamics on initial conditions. For a randomly chosen initial state, the dynamics will be indistinguishable from that of a generic quantum ergodic system, and the expectation values of local observables will quickly decay to the ensemble average. However, when prepared in specific states whose spectral decomposition is dominated by the scarred eigenstates, very different behaviour is observed. For example, the Loschmidt echo $\mathcal{L}(t) = |\langle \psi(t) | \psi(0) \rangle|^2$, which in an ergodic system should decay quickly to (almost) zero and stay there as the dynamics explore the Hilbert space, instead periodically




Figure 3.8: Infinite-temperature autocorrelation functions (3.21) at the edge and in the bulk for (a) Z_ℓ and (b) K_ℓ . Data is obtained by exact diagonalisation for different system sizes L as indicated by the arrows. Note that data in (a) is analogous to Ref. [28]. The dashed lines in (b) signal the equipartition value $1/(L-2)$. The other parameters are chosen as $V = 0.05$, $\Gamma = 0.1$, and $\lambda = 0.6$.

returns to a large value at later times or decays extremely slowly. Local observables likewise may depart far from equilibrium values, but again only for certain initial states. Additionally, Ref. [28] showed that the SPT edge modes in the ZXZ model were protected by dimerisation, but did not find evidence of preservation of bulk cluster operators. We therefore considered whether the presence of scars, and a prudent choice of initial state, would in fact produce anomalous dynamics in the bulk. Particularly, as the unique feature linking the non-thermal states is that one or both sublattices are fully polarised, we focus on cluster basis states (see Sec. 3.4.1) with this property.

Firstly, consider the infinite-temperature autocorrelation functions of local operators A_ℓ ,

$$\langle A_\ell(t)A_\ell \rangle_\infty = \frac{\text{Tr}[A_\ell(t)A_\ell]}{2^L}, \quad (3.21)$$

where $A_\ell(t) = e^{+i\mathcal{H}t}A_\ell e^{-i\mathcal{H}t}$, for both Pauli spin ($A_\ell = \sigma_\ell^z$) and cluster ($A_\ell = K_\ell$) operators. In the first case, Fig. 3.8(a), there is a clear difference in behaviour between the edge spin and the bulk, with a rapid decay for $\ell = L/2$ but a lifetime exponentially long in system size for $\ell = 1$. This is in fact the main result from Ref. [28]. For the cluster operators in Fig. 3.8(b), looking at the bulk ($\ell = L/2$), the autocorrelation function again decays rapidly, but strangely does not reach either zero or the equipartition value $1/(L-2)$ which would be given if the cluster




Figure 3.10: Averaged cluster expectation value $\langle K_{e/o}(t) \rangle_{|\psi_2\rangle}$ on even/odd sublattice [Eq. (3.23)] for the initial state $|\psi_2\rangle$ [see Figs. 3.9(b),(e)]. Data is shown for system sizes $L = 14, 16, 18, 20$. In panels (a) and (b), we show data for the dimerised model with $\lambda = 0.6$, while panel (c) shows data for $\lambda = 1$. In (a) and (c), we have $\Gamma = 0.1$ and $V = 0.05$. Panel (b) considers slightly stronger perturbations $\Gamma = 0.2$ and $V = 0.1$.

In Fig. 3.9, we elucidate the effect of λ on the quench dynamics $\langle K_e(t) \rangle_{|\psi\rangle}$ for three exemplary initial states $|\psi_1\rangle, |\psi_2\rangle, |\psi_3\rangle$,

$$\langle K_\ell(t) \rangle_{|\psi\rangle} = \langle \psi(t) | K_\ell | \psi(t) \rangle, \quad (3.22)$$

with $\Gamma = 0.1$, and $V = 0.05$ in all cases. Panels (a)-(c) have $\lambda = 0.6$, which is in the dimerised regime in which scars were detected and Ref. [28] found enhanced coherence of the edge modes, while panels (d)-(f) have $\lambda = 1$. For $|\psi_1\rangle$ and $|\psi_2\rangle$, the odd sublattice is fully in the ground state, and so these states fall into the scarred manifold in the dimerised model. In fact, $|\psi_2\rangle$ can be exactly identified with the approximation to the highest energy state in the tower, $|T_{L/2-1}\rangle$. However, $|\psi_3\rangle$ has cluster excitations on both sublattices, and we would therefore expect it to thermalise rapidly regardless of the value of λ . This is exactly what we observe: in Fig. 3.9(a) and (b), we observe that excitations remain confined to their respective legs and the even sublattice remains in the ground state. However, in (c), the excitations spread uniformly throughout the system and the end result is a thermal state. When we leave the dimerised regime in panels (d)-(f), all three states thermalise rapidly. This difference in dynamical behaviour was already shown in Fig. 3.2(d) through the Loschmidt echo $\mathcal{L}(t)$, which exhibited clear revivals for a charge density wave state, but not a different cluster state with both odd and even excitations.

Additionally, in Fig. 3.10, we show for $|\psi_2\rangle = |T_{L/2-1}\rangle$ the average cluster




Figure 3.11: Overlap of the states $|\psi_1\rangle$ – $|\psi_3\rangle$ (c.f. Fig. 3.9) with the eigenstates $|n\rangle$ of $\mathcal{H}$. The insets show the entanglement entropy S_A versus subsystem size for the eigenstate $|n\rangle$ which has the largest overlap (indicated by the blue asterisk). The parameters are chosen as $\Gamma = 0.1$, $V = 0.05$, and $\lambda = 0.6$, and the system size is $L = 14$.

operator expectation $\langle K_{e/o}(t) \rangle_{|\psi\rangle}$ on the even (e) and odd (o) sublattices,

$$\langle K_e(t) \rangle_{|\psi\rangle} = \frac{2}{L-2} \sum_{\ell=1}^{L/2-1} \langle \psi(t) | K_{2\ell} | \psi(t) \rangle, \quad (3.23)$$

and with $K_{2\ell} \rightarrow K_{2\ell+1}$ in the case of $\langle K_o(t) \rangle_{|\psi\rangle}$. We show this data in three regimes: in both of (a) and (b) we set $\lambda = 0.6$, and either $\Gamma = 0.1$ and $V = 0.05$ in (a), or double these values in (b). Panel (a) shows that the cluster excitations are stable for times at least up to $t \simeq 180$, and in fact we have further numerical evidence (not shown here) that the values remain far from zero up to $t > 10^5$. Panel (b) repeats this numerical experiment in a regime in which the level statistics are clearly Wigner-Dyson (see Sec. 3.3.4), in which case we still see preservation of the cluster excitations although the decay is quicker (c.f. Fig. 3.9). However, for $\lambda = 1$ in panel (c), the values decay rapidly to zero by $t \simeq 20$.

It is hence clear from Figures 3.9 and 3.10 that dimerising the model enables dynamics with long coherence times not only at the edges like in Ref. [28], but also in the bulk when choosing appropriate operators and initial states.

Furthermore, the spectral decomposition of the states $|\psi_{1,2,3}\rangle$ in Fig. 3.11 shows that the first two are dominated by a small number of energy eigenstates at the same energy, which appear to have sub-volume law entanglement entropies. However, $|\psi_3\rangle$ contains contributions from a large number of eigenstates, each with a small coefficient, and these are clearly volume law. This provides a clear link between the




Figure 3.12: (a) Time- and space-averaged correlation function $\overline{K_{|\psi\rangle}}(t)$ (3.24) and (b) Loschmidt echo $\mathcal{L}(t) = |\langle\psi(t)|\psi\rangle|^2$, for the dimerised model with $\lambda = 0.6$ and fixed system size $L = 12$. Data is shown those states $|\psi\rangle$ from the cluster basis with an energy $E_{|\psi\rangle} \in [-0.5, 0.5]$ in the centre of the spectrum. Initial states where one of the two sublattices is fully excited/fully in the ground state (red curves) saturate to high values of $\overline{K_{|\psi\rangle}}(t)$ and exhibit distinct oscillations in $\mathcal{L}(t)$.

existence of non-thermal sub-volume law entangled eigenstates in the spectrum of $\mathcal{H}$, and the stable nature of certain initial states which overlap heavily with these eigenstates.

Finally, we studied quench dynamics from the wide class of cluster basis states, using the time- and space-averaged correlation function $\overline{K_{|\psi\rangle}}(t)$ (similar to that used in Ref. [213]),

$$\overline{K_{|\psi\rangle}}(t) = \frac{1}{L-2} \sum_{\ell=2}^{L-1} \frac{1}{t} \int_0^t \langle\psi|K_\ell(\tau)K_\ell(0)|\psi\rangle d\tau . \quad (3.24)$$

If the long-time value of $\overline{K_{|\psi\rangle}}(t)$ is non-zero or even close to one, this would indicate that memory of initial conditions, specifically that of local cluster operators, is preserved. Here we consider only those states with an energy close to the middle of the spectrum,

$$-0.5 \leq E_{|\psi\rangle} \leq 0.5 , \quad E_{|\psi\rangle} = \langle\psi|\mathcal{H}|\psi\rangle . \quad (3.25)$$

If the system obeys the strong ETH, then initial states at roughly the same energy would yield very similar values of local observables at long times. Yet Fig. 3.12(a) shows that actually, for a certain class of states in the dimerised model, this is violated: while the correlator $\overline{K_{|\psi\rangle}}(t)$ decays quickly to zero for the majority of initial states, for a small number of states the long-time value is clearly non-zero.

And in fact, this class of states is exactly those with a fully polarised sublattice, identified earlier as the non-thermal manifold. Focusing also on the Loschmidt echo $\mathcal{L} = |\langle \psi(t) | \psi(0) \rangle|^2$ in Fig. 3.12(b), we see that while for all states the echo decays rapidly to zero, for certain states in which one sublattice is polarised the data exhibits distinct revivals. These revivals do not occur when $\lambda = 1$. We argue that this is the direct result of the decoupling of the two sublattices under $\lambda \neq 1$, which in turn leads to an effective fragmentation of the Hilbert space into weakly connected subspaces. Specifically, those subspaces where one sublattice is polarised are weakly coupled to the rest of the Hilbert space, and so initial states in these manifolds initially fail to explore the full Hilbert space under time evolution such that some local memory of initial conditions is preserved. This is strong evidence, therefore, of the existence of quantum scars in the ZXZ model.

3.3.6 Discussion

In this section, we have studied the eigenstate properties and out-of-equilibrium dynamics of an interacting spin-1/2 chain with three-site “cluster” terms, which is known to host a $\mathbb{Z}_2 \times \mathbb{Z}_2$ symmetry-protected topological phase. In particular, we worked in a regime with weak interactions where the dimerisation parameter $\lambda \neq 1$, such that cluster excitations on the odd and even sublattices respectively are energetically detuned, and which Ref. [28] found stabilises a long-lived boundary mode even at infinite temperature.

In this regime, we uncovered a set of atypical states throughout the spectrum which do not obey the ETH ansatz, and which we identify as quantum many-body scars. These scars form a number of “towers of states”, connected by raising and lowering operators, and approximated by a simple ansatz expressible in matrix-product state form. Looking at the properties of the scars, we find that their key signature is a fully polarised sublattice with regards to cluster excitations, such that the even and odd sublattices have very different filling fractions. Additionally, they have abnormally high overlaps with certain low-entanglement states which might be prepared experimentally, and sub-volume law entanglement entropies. Most intriguingly, there is evidence that these scar states preserve signatures of SPT order, namely a four-fold degenerate entanglement spectrum.

With the parameters chosen, the ZXZ model is not strongly chaotic, but in fact

in an intermediate regime between integrable and chaotic level statistics. This is different to the phenomenology of other scar models, which are typically fully ergodic with Wigner-Dyson level statistics. On the one hand, this may be due to the scarred manifold occupying a relatively large fraction of the total Hilbert space at the system sizes accessible to numerics, compared to in these other models. On the other hand, this may be unconnected to the scars, for example being due to the presence of edge modes or the relatively weak interaction term.

We also looked at the out-of-equilibrium dynamics of the model, establishing the importance of the dimerisation parameter λ and the effect of scars on the dynamics. With $\lambda = 1$, cluster excitations are free to spread out over the entire system, which rapidly thermalises to a uniform state no matter the initial configuration. However, with $\lambda = 0.6$ as in Ref. [28], we observe that the cluster excitations remain confined to the odd or even sublattices, weakly fracturing the Hilbert space and preventing full thermalisation. This effect is also seen to preserve or lead to the revivals of certain dynamical signatures, such as correlation functions or the Loschmidt echo, when the initial state has a fully-polarised sublattice, which shows that the atypical dynamics may be due to the quantum scars discovered in this model. We note that these discoveries do not rely on the presence of edge modes – we have observed similar phenomena to that in Fig. 3.9 with periodic boundary conditions.

This work raises a number of questions. While the ZXZ model in the regime we have investigated is certainly not integrable, there is evidence that it is not fully chaotic either. Furthermore, we found that the scars we discovered were also present in the integrable regime with $V = 0$. Therefore, we consider whether it may be worth extending the notion of scars to integrable models. While these do not thermalise in the traditional sense, they may be described by the generalised Gibbs ensemble (GGE) which takes account of the extensively many local conservation laws [214, 215], and so it is reasonable to wonder whether scars might also interfere with thermalisation to the GGE. Noting also that the non-interacting ZXZ model is equivalent to two copies of the transverse-field Ising model (TFIM), it would be interesting to explore connections between these scars and the well-known local conservation laws in the TFIM with open boundary conditions, as well as the stability of these laws to perturbations. We note that the connection between scars and proximity to an

integral point has already been discussed in Refs. [160, 216].

The role of disorder on the topological nature of quantum scars in the ZXZ model is a question of interest in the context of many-body localisation. It has been shown for example that the PXP model can be localised [167], but the constraints in the model also induce interactions which may prevent MBL in the thermodynamic limit [168]. Moreover, relatively weak disorder may destroy scars [166], even though at strong disorder many-body localisation reinstates the topological edge modes. Studying the phase transition between various finite-temperature phases with topological features, in intermediate regimes between scars and localisation, could therefore provide new insights into all three of these phenomena, and the relationships between them. On the topic of topological order, one could also consider searching for (naturally arising) scars in higher-dimensional systems which preserve full topological order, as opposed to its weaker symmetry-protected cousin. The ZXZ model also has higher-dimensional extensions [196, 200, 201] – it is unknown at present if such systems host a scarred manifold. Ref. [28] also notes that the ZXZ chain can be realised through Floquet engineering applied to a 1D optical-tweezer array, raising the possibility that the unusual dynamics observed in this work could be verified in experiment.

Finally, we note that while single particle quantum scars have typically been connected to an unstable periodic orbit in a corresponding classical system [20, 147–150] (see Sec. 2.3.3), there is no such obvious connection for quantum many-body scars. However, recent work applying the time-dependent variational principle (TDVP), a method for constructing semiclassical trajectories through a variational manifold which best approximate dynamics in the full Hilbert space [170, 171, 217], to the PXP model has uncovered evidence of exactly these unstable orbits [20, 173, 218]. These orbits are able to explain the oscillatory dynamics observed in Rydberg atom simulators in Ref. [9], as well as quantify their stability, providing for the first time a clear connection between quantum scars and semiclassical dynamics. This research was expanded upon later to also uncover such semiclassical trajectories corresponding to a $|\mathbb{Z}_3\rangle$ initial state, as well as applying the method to a 2D model using a tensor tree state ansatz [218]. We could therefore consider applying this same method to the ZXZ model, utilising a variational ansatz based on the MPS description in Appendix 3.4.2; a similar construction was used in Ref. [173], whereby a spin product

state was projected into the subspace without adjacent excitations. However, there is no guarantee this is the correct choice of ansatz.

Since the research in this chapter was performed, there have been many further developments in the field of QMBS, including some involving the ZXZ model. In Ref [219], it is shown that a certain version of the ZXZ model, which has spatial inversion symmetry but anticommutes with the spin parity operator $\prod_j \sigma_j^z$, acquires an exponentially degenerate nullspace containing low-entanglement eigenstates. This property is not specific to the ZXZ model, but instead arises through a generic mechanism dependent only on the symmetry considerations above. While small modifications to the Hamiltonian (3.5) would cause it satisfy these conditions, it is not clear if the scars would survive these changes. Additionally, rather than using TDVP to explain known examples of quantum scars, Ref. [174] has used the method to engineer families of Floquet models which contain desired scarred orbits. Specifically, this work was able to enhance the stability of the orbits by suppressing leakage from the TDVP manifold. While that work considers a driven AKLT model, the technique can also be applied to a driven ZXZ model [220]. We also note that Ref. [221] suggests that, in models whose classical limits exhibit certain kinds of frustration, SPT order could in fact *lead* to quantum scars (although this is not demonstrated).

3.4 Supplementary Material

3.4.1 Exact eigenstates of the clean ZXZ model

There exists a set of states $|\psi_{\{\pm\ell\}}^{g_e, g_o}\rangle$ which are mutual eigenstates of the cluster operators and have definite $\mathbb{Z}_2 \times \mathbb{Z}_2$ symmetry values, forming a complete basis within each symmetry sector.

Let us define the state $|\psi_{\{\pm\ell\}}^{g_e, g_o}\rangle$ such that,

$$\hat{G}_{e/o} |\psi_{\{\pm\ell\}}^{g_e, g_o}\rangle = g_{e/o} |\psi_{\{\pm\ell\}}^{g_e, g_o}\rangle, \quad (3.26)$$

$$K_\ell |\psi_{\{\pm\ell\}}^{g_e, g_o}\rangle = \pm_\ell |\psi_{\{\pm\ell\}}^{g_e, g_o}\rangle, \quad 2 \leq \ell \leq L-1, \quad (3.27)$$

where $\hat{G}_{e/o}$ are the symmetry operators for even and odd site spin-flip symmetry respectively, and $g_{e/o} = \pm 1$ their eigenvalues. This does not uniquely specify the

phase of these states, but they can be defined concretely for our purposes as,

$$|\psi_{\{\pm\ell\}}^{g_e, g_o}\rangle = \frac{1}{\sqrt{2^{L-2}}} \left(\prod_{j=2}^{L-1} K_\ell^{\pm\ell} \right) P_e^{g'_e} P_o^{g'_o} |\uparrow\rangle^{\otimes L}, \quad (3.28)$$

where $P_{e/o}^\pm = (1 \pm \hat{G}_{e/o})/2$, and $g'_{e/o} = g_{e/o}^{L/2-1} = \pm 1$. The state $|\psi_{\bar{1}}/0\rangle$ is constructed by applying lowering operators only on the even sites, and adjusting the normalisation accordingly.

It is easy to check that these states have the correct symmetry eigenvalues, after noting that $\{K_{2j}^\pm, G_e\} = [K_{2j}^\pm, G_o] = 0$ and likewise $[K_{2j+1}^\pm, G_e] = \{K_{2j+1}^\pm, G_o\} = 0$. Noting also that $K_\ell K_\ell^\pm = \pm K_j^\pm$ shows it has the correct cluster-operator eigenvalues.

It is simple to show that these states are mutually orthogonal. Take two states which differ in their cluster eigenvalues at at least one site. Since $(K^\pm)^\dagger = K^\mp$, taking the inner product of these states, substituting in (3.28) and grouping the operators by site (using the fact that operators on different sites commute) will lead to at least one factor $(K^\pm)^2 = 0$. Hence the inner product must vanish. It is clear also that states with different symmetry eigenvalues must be orthogonal.

Since there are $L - 2$ sites each with two choices for cluster eigenvalues, and also four symmetry sectors, there are $4 \times 2^{L-2} = 2^L$ cluster states – exactly the number of states in a system of size L . Hence this is a complete, orthonormal basis.

3.4.2 Matrix Product State representation

In this section we show that the states $|T_j(k)\rangle$ in fact admit a matrix product state (MPS) representation with a maximum bond dimension linear in L . This is a representation of a quantum state in the form,

$$|\psi\rangle = \sum_{\{\sigma_i\}} M_1^{\sigma_1} M_2^{\sigma_2} \dots M_N^{\sigma_N} |\sigma_1 \sigma_2 \dots \sigma_N\rangle, \quad (3.29)$$

where the M_i are rank-3 tensors with a “physical” index σ_i and two “auxiliary” indexes of dimension χ (suppressed in the above equation), typically represented using a tensor network [222]. These are usually very efficient representations of low-entanglement states, and in particular translationally-invariant area-law states can be represented with a finite number of elements. Conversely, an MPS of constant bond dimension χ is guaranteed to be area law, as the Schmidt decomposition of




Figure 3.13: Matrix product state representation of a tower state $|T_j(k)\rangle$. The tensors A, B generate j delocalised excitations with momentum k on the odd sublattice, and the tensors C map this from a cluster basis description to the physical spin basis. When the vertical edges K_ℓ are contracted, we obtain an MPS of maximal bond dimension $\chi = 4[\max(j+1, L/2-j)] \leq L+4$. Triangles indicate legs which are held constant (achieved by contracting with boundary vectors A_l and A_r).

some partition can be obtained directly from the MPS and the number of coefficients is equal to χ .

We do this in two steps; see Fig. 3.13 for reference. First, we construct j cluster-wave excitations of momentum k on the odd sublattice (tensors A, B , in blue), and then we map that state from the cluster basis to the spin basis via a matrix-product operator (MPO) of dimension $\chi = 4$ (tensors C , in green).

We demonstrate this construction first for the simple case $j = 1$, i.e. a single delocalised cluster excitation on the odd sublattice. We can write this state in the cluster basis (App. 3.4.1) as,

$$|\psi\rangle = \sum_{\ell=1}^{L/2-1} e^{ik\ell} |\bar{1}\bar{1} \cdots 1_{2\ell+1} \cdots \bar{1}\bar{1}\rangle. \quad (3.30)$$

We can then write this in matrix product form as,

$$|\psi\rangle = \sum_{\{K_\ell\}} A_l \left(\prod_{\ell=1}^{L/2-1} B_{K_{2\ell}} A_{K_{2\ell+1}} \right) A_r |K_2 K_3 \cdots K_{L-1}\rangle, \quad (3.31)$$

$$A_{+1} = \begin{pmatrix} 0 & 0 \\ 1 & 0 \end{pmatrix}, \quad A_{-1} = \begin{pmatrix} 1 & 0 \\ 0 & \alpha \end{pmatrix}, \quad (3.32)$$

where $A_l = (0, 1)$ and $A_r = (1, 0)^T$, $\alpha = e^{ik}$, and $B_K = \mathbb{I} \delta_{K, -1}$ selects states with $K_{2\ell} = -1$ on even sites.

Consider a single element of the summation in Eq. (3.31). Note that $A_{-1}^n A_r = A_r$ but $A_{-1}^n A_{+1} A_r = \alpha^n A_l^T$, and $A_{+1}^2 = 0$. Hence it is clear there must be exactly one

A_{+1} in the matrix product or else it will vanish, while we accumulate a factor α for each site to the left of this A_{+1} . This gives the desired state, up to a constant phase factor which we can ignore.

In the above example, the index n of the nonzero entry in the vector as we moved from right to left counted the number of excitations to the right of the current site. To extend this idea to j such excitations, we need a bond dimension of $j + 1$ such that $0 \leq n \leq j$,

$$A_{+1} = \begin{pmatrix} 0 & & & & \\ 1 & 0 & & & \\ & \alpha & \ddots & & \\ & & \ddots & 0 & \\ & & & \alpha^{j-1} & 0 \end{pmatrix}, \quad A_{-1} = \begin{pmatrix} 1 & & & & \\ & \alpha & & & \\ & & \alpha^2 & & \\ & & & \ddots & \\ & & & & \alpha^j \end{pmatrix}, \quad (3.33)$$

where we now take $A_l = (1, 0, \dots, 0)$ and $A_r = (0, \dots, 0, 1)^T$. Note that now a factor α^n is accumulated at each site when there are n excitations to the right.

We now transform this state $|\psi\rangle = |T_j(k)\rangle$ to the physical spin basis using an MPO formed of the rank-4 tensors $C_\sigma^K = \text{diag}(c_\sigma^K, c_\sigma^K)$ where,

$$c_\downarrow^K = \begin{pmatrix} 0 & K \\ 0 & -K \end{pmatrix}, \quad c_\uparrow^K = \begin{pmatrix} 1 & 0 \\ 1 & 0 \end{pmatrix}. \quad (3.34)$$

The tensors $c_\downarrow^K$ here form the known MPS description for a cluster state, where $K = \pm 1$ is the eigenvalue of K_ℓ for the state at each site [222]. Finally, the tensors C_l and C_r are the contraction of the left or right leg respectively of C_σ^{+1} with an appropriate boundary vector – this determines the symmetry sector. This MPO hence has bond dimension $\chi_C = 4$

When the vertical edges representing K_ℓ are contracted, the resultant MPS will have maximal bond dimension $\chi = \chi_A \chi_C = 4(j + 1)$, which is $O(L)$ since $j < L/2$. Certain optimisations can improve this to $4 \max(j + 1, L/2 - j)$. This bounds the entropy growth of the $|T_j(k)\rangle$ by $O(\log L)$.




Figure 3.14: Overlaps of the states $|S_j^\pm(k)\rangle$ and $|T_j^\pm(k)\rangle$, in analogy with the inset of Fig. 3.3, for states constructed with a generalised excitation operator $\mathcal{O}(k)$ (3.35). Data for the non-interacting ($V = 0$) model is in blue, and for the interacting model ($V = 0.05$) in orange. [Other parameters: $L = 16$, $\lambda = 0.6$, $\Gamma = 0.1$, $G_e = G_e = +1$]

3.4.3 Non-zero momentum excitations

We can generalise Eq. (3.16) to create excitations with an arbitrary momentum k simply by introducing a location-dependent phase,

$$\mathcal{O}(k) = \sum_{\ell=1}^{L/2-1} e^{ik\ell} K_{2\ell+1}^+ . \quad (3.35)$$

Because the Hamiltonian is real symmetric, the eigenstates can be taken to be real. However, the towers $|T_j(k)\rangle$ formed from these operators are in general complex, so we must take linear combinations to produce real states. Since $|T_j(k)\rangle$ and $|T_j(-k)\rangle$ are complex conjugates, we can consider the states,

$$|T_j^\pm(k)\rangle = \frac{1}{\sqrt{2}} (|T_j(k)\rangle \pm |T_j(-k)\rangle) , \quad (3.36)$$

which are either real or have a global phase which can be eliminated by multiplication by a constant. We can also define $|S_j^\pm(k)\rangle$ in analogy with the definition in the main text. Note that $|T_0(k)\rangle = |\text{gs}\rangle$, and $|T_{L/2-1}(k)\rangle$ is the state with a fully-excited odd sublattice and ground-state even sublattice, for all values of k , and so for these states we do not take linear combinations.

Figure 3.14 provides some data on the squared overlap, $|\langle T_j^\pm(k)|S_j^\pm(k)\rangle|^2$, between these states and the eigenstates of the model. As an indicative example, in the first two panels we show data for $\eta = 2, 3$, where $k = 2\pi\eta/(L/2 - 1)$. For

$\eta = 2$, the first excited state ($j = 1$) and the penultimate state ($j = L/2 - 1$) are close approximations, regardless of interactions, however the states in the middle of the tower are significantly worse. For $\eta = 3$, all the states are close approximations to eigenstates in the non-interacting model, and remain the majority spectral contribution when interactions are turned on. The overlaps for other values of η generally follow a similar pattern to one of these two cases.

The final panel shows that half-integer values of η can produce good overlaps too; particularly, we found that for $\eta = 1/2$, the overlap for the first excited state is very high, and even better than for $k = 0$ in Sec. [3.3.3.2](#).

Chapter 4

Resonance proliferation between many-body localised phases

4.1 Introduction

For generic interacting quantum systems, time evolution scrambles quantum information on the approach to a thermal state, which limits our ability to efficiently process that information. The eigenstates of thermalising systems satisfy the eigenstate thermalisation hypothesis (ETH) [3, 5, 49, 223] and encode information in exponentially complex observables, rendering them irretrievable (see Sec. 2.1). Exceptions to the ETH are therefore important for the encoding and manipulation of information in weakly entangled quantum states; likewise, attempts to control and manipulate quantum systems for processing quantum information have inspired great advances in our understanding of quantum dynamics. In clean systems, the prototypical examples are the integrable systems, which possess an extensive number of integrals of motion [46, 47], and scarred many-body Hamiltonians with a manifold of ETH-violating eigenstates [9, 10, 18]. These both require an element of fine-tuning which leaves them unstable to perturbation. Instead, many-body localisation (MBL) [7, 8, 36, 67, 75, 90, 93, 97, 98] (see Sec. 2.2) provides a promising avenue towards *robustly* avoiding thermalisation. The emergence of local integrals of motion (l-bits) in the MBL phase [101, 108, 113, 116, 117, 224, 225] allows for the effective protection of *classical* information but, as explained in Sec. 2.2.3, this performs poorly with quantum information [126, 127].

The situation at zero temperature is a little different, where topological order and symmetry-protected topology (SPT) allow for the robust encoding of quantum

information into a degenerate ground-state manifold [27] (see Sec. 2.4). These are stable to (symmetry preserving) perturbations, provided the energy gap to the excited states does not close, but typically fail in the presence of thermal noise and delocalised excitations [24, 176, 184] as the degeneracy in the ground states is usually not replicated in the highly-excited states. The addition of disorder to a system with ground-state SPT order provides an avenue to stabilise that order at finite energy densities [29–31, 35, 178] – producing a system with multiple topologically distinct MBL phases and possibly direct eigenstate-ordering phase transitions between them. Additionally, even in clean systems the presence of a strong zero mode (SZM) imposes a spectrum-wide energy pairing [128, 186, 226] and can enable coherent storage of quantum information [28], as per Sec. 2.4.2. However, recent exact diagonalisation studies demonstrate that an ergodic phase may intervene at arbitrarily small interaction strengths, and some even claim that an MBL-to-MBL phase transition is forbidden [32–35]. Avalanches induced by rare regions and resonances [37–45] play a crucial role in destabilising MBL at the localisation transition, and are candidates for generating the intervening delocalised phase here (see Sec. 2.2.4).

Finite-depth tensor network techniques [110–114], flow-equations [107–109], and renormalisation group (RG) approaches [103–105] for approximating weakly entangled excited states provide access to dynamical properties, critical behaviour, and l-bit operators, and have enriched our understanding of MBL. The RG techniques progressively eliminate or “decimate” degrees of freedom from a system, typically starting with the smallest length scales and highest frequencies, to arrive at a long-range or low-energy effective model. For example, one may integrate out the highest-momentum excitations while preserving long-distance correlations, or coarse grain in real space by grouping local degrees of freedom (such as spins). At each stage, the Hamiltonian is then renormalised to best recreate the physics prior to rescaling, with the change in parameters known as the RG flow. It is generally then hoped that the RG flow will lead to certain terms in the Hamiltonian dominating while others, known as irrelevant operators, vanish, with the flow terminating at certain stable fixed points in parameter space. These fixed points, and all those regions of parameter space flowing into them, are then identified with phases of

matter whose physics is controlled by the properties of the fixed point.

For disordered systems, we may make use of the real space RG for excited states (RSRG-X) [31, 106]. In the traditional ground-state RSRG (also known in the literature as the strong disorder RG, or SDRG), one locates the local term in the Hamiltonian responsible for the largest gap in the system and then solves for the ground-state manifold, removing one or more degrees of freedom in the process [227–229]. The Hamiltonian is then replaced by the effective Hamiltonian on this manifold, and the process repeats until a low-energy description of the system is obtained. RSRG-X extends this by instead allowing one to pick either the ground-state or excited manifold at each step, hence arriving at an effective description of a state at finite energy density. The coarse-graining process lends itself to identifying the real-space structure of resonances, and studying their size and statistics; indeed, this has been looked at for a certain model with $SU(2)$ symmetry, which exhibits an MBL-like regime a finite times and finite sizes [151, 152], showing also that RSRG-X can be used to model states that are not in fact strictly localised.

In this work we have applied RSRG-X to an interacting spin-1/2 chain with two MBL phases: a trivial paramagnetic phase and a spin glass phase with SPT order protected by a global $\mathbb{Z}_2$ symmetry. We have extracted a Clifford circuit and Schrieffer-Wolff transformation which together approximately diagonalise the Hamiltonian to first order [106]. These encode the localised basis that would best fit the eigenstates of the Hamiltonian *if* the system were localised, and we probe its stability to the off-diagonal part of the Hamiltonian by searching for many-body resonances [135, 151, 152, 230] between these basis states. Additionally, the geometry of these resonances, and in particular how these link l-bits together into thermal clusters, allows us to investigate the breakdown of localisation through the use of a finite size scaling analysis. While Ref. [33] looked at the size of resonances in this same model, this is the first study to look at the structure of resonances in this manner.

We find that the marginal MBL phase, as found in Ref. [31], is indeed destabilised to an ergodic phase for even relatively small interaction strengths, and that this phase may be extended in parameter space even with infinitesimal interactions. We show that the resonances filter through the system to form clusters that scale extensively

with system size in the ergodic phase. We also look at the variance of the energy δH^2 of the RSRG-X basis, which quantifies the accuracy of these states as approximations to the true eigenstates.

The rest of this chapter proceeds as follows. In the next section, we lay out the details of the interacting Ising-Majorana model, and develop the Clifford RSRG-X technique as well as its application to this model. Then, in Sec. 4.3, we present the results of this work, including: the discovery of a strong edge zero mode in the localised SPT phase in Sec. 4.3.1; the calculation of energy variance in Sec. 4.3.2; a description of resonant mixing in the RSRG-X basis in Sec. 4.3.3; the spatial distribution of these resonances in Sec. 4.3.4; and finally the scaling with increasing system size in Sec. 4.3.5, a key result of this work. We finish in Sec. 4.4 with a discussion of these results, giving our conclusions and suggesting future avenues of research.

4.2 Model and RSRG-X


We consider a transverse field Ising model with nearest-neighbour and next-nearest-neighbour interactions, also known as the interacting Ising-Majorana model, described by the Hamiltonian

$$H_{\text{IM}} = \sum_i h_i \sigma_i^z + J_i \sigma_i^x \sigma_{i+1}^x + g (\sigma_i^z \sigma_{i+1}^z + \sigma_i^x \sigma_{i+2}^x) , \quad (4.1)$$

with $h_i \sim \text{Uniform}[0, h]$ and $J_i \sim \text{Uniform}[0, J]$, which we normalise by setting $hJ = 1$. We also use open boundary conditions. This model is statistically self-dual under the exchange $h \leftrightarrow J$, and is known to have two distinct MBL phases [29, 31, 178]. On one hand, when h is large, the local fields dominate and the model enters a topologically trivial paramagnetic (PM) phase, where the energy eigenstates are products states of frozen spins aligned along the z -axis. On the other hand, when J is large, the model enters a spin glass (SG) phase, with spins forming large entangled clusters due to the action of the nearest-neighbour $\sigma_i^x \sigma_{i+1}^x$ terms. In this phase, the system is topologically ordered, protected by the global parity symmetry $G = \prod \sigma_i^z$, and hosts a Majorana edge zero mode [186, 189]. This zero mode leads to spectral pairing between the two parity sectors, with the gap exponentially small in system size [32]. We characterise the phase of the model by

the quantity $\delta = \overline{\ln|J_i|} - \overline{\ln|h_i|} = 2 \ln J$, such that the model is dual about $\delta = 0$ with positive and negative delta in the SG and PM phases respectively.

In fact, if we choose to represent this Hamiltonian (via the Jordan-Wigner transformation) in terms of Majorana fermions such that two such fermions $(\gamma_{2i}, \gamma_{2i+1})$ represent each physical spin σ_i , in an infinite or periodic chain the statistical duality above is made exact by regrouping the fermions into a new set of spins $\tilde{\sigma}_i$ each represented by $(\gamma_{2i-1}, \gamma_{2i})$:



Ground state RSRG has been used extensively to characterise similar models, including the model (4.1) in the fermionic representation [231], but here we are interested in the excited states. Applied rigorously, RSRG-X relies on an assumption of strong disorder, with relevant couplings distributed according to some power law $P(x) \propto x^\alpha$ (with $\alpha > 1$), leading to a good separation of energy scales in the system. In many cases the assumption of strong disorder may be relaxed: the system will quickly flow towards the infinite-randomness fixed point, and so the validity of the procedure is preserved. For all work in this chapter, we set $g \ll \max(h, J)$, ensuring that one of the couplings h_i and J_i is always the largest in the system and thus the only couplings that need to be directly considered by the RSRG-X procedure. This keeps the Hamiltonian to a closed form. While the procedure can in principle generate a dominant interaction coupling, this is rare so long as g is not too large, and we assume that this occurs infrequently enough not to meaningfully affect the disorder-averaged data.

We therefore consider two types of decimation: the freezing of a single spin due to a dominant local field $h_i \sigma_i^z$ (“site decimation”) and the merger of two spins into one due to a dominant bond $J_i \sigma_i^x \sigma_{i+1}^x$ (“bond decimation”). In each case, the local Hilbert space is first rotated via an approximate Schrieffer–Wolff (SW) transformation [232], truncated at second order, towards a basis aligned with the gap [31, 106], and then projected onto the subspace above or below this gap where the (transformed) local operator corresponding to the leading term (σ_i^z or $\sigma_i^x \sigma_{i+1}^x$) is equal to $c = \pm 1$. The

perturbative transformation removes terms that anticommute with the leading term, but also produces new terms which are second order in sub-leading energy scales. These new terms physically originate from the combination of two removed terms, mediated by the decimated spin: for example, when the term $h_2\sigma_2^z$ is decimated, two nearest-neighbour terms $J_1\sigma_1^x\sigma_2^x$ and $J_2\sigma_2^x\sigma_3^x$ (which both anticommute with $h_2\sigma_2^z$) are combined to form a term $c(J_1J_2/h_3)\sigma_1^x\sigma_3^x$. At higher orders there can be contributions producing the same terms as those produced at second order and these interfering contributions are not included. The full RG rules may be found in Appendix 4.5.1.

We will later incorporate the SW transformations into our ansatz but if we neglect these, the successive merger of spins due to bond decimations causes the RG states to acquire a tree-like structure. As such these states can in fact be represented by tree tensor networks (TTNs) [151, 152, 233] where each node in the network represents a decimation, with n incoming legs of bond dimension $d = 2$ for the spins to be decimated, and $n - 1$ outgoing legs for the new effective spins after the decimation. The network takes us from L physical spins and successively removes each degree of freedom, narrowing with each RG step. The tensor for a bond dimension is a projector from two spins onto one of the $\sigma_i^x\sigma_{i+1}^x = \pm 1$ subspaces, an isometry with two incoming legs and one outgoing leg. To represent spin up and down in the renormalised σ^z basis, we choose respectively:

$$|a_c\rangle = \frac{1}{\sqrt{2}}(|\uparrow\uparrow\rangle + c|\downarrow\downarrow\rangle) , \quad (4.3)$$

$$|b_c\rangle = \frac{1}{\sqrt{2}}(|\uparrow\downarrow\rangle - c|\downarrow\uparrow\rangle) , \quad (4.4)$$

with $c = \pm 1$, such that the isometry is given by $|\uparrow\rangle\langle a_c| + |\downarrow\rangle\langle b_c|$. The site decimation tensors are then projections of a single spin onto $\sigma_i^z = \pm 1$ – just the two basis vectors in $d = 2$. This is represented pictorially in column (i) of Fig. 4.1, where we show typical tree tensor networks for three points in the phase diagram: the spin glass phase, the paramagnetic phase, and the critical phase at $\delta=0$.

At each step we must choose between the excited and ground state manifolds, by selecting $c_k = \pm 1$. Since the energy shift of each decimation step typically decreases throughout the procedure, producing a hierarchy of energy scales, we can view the




Figure 4.1: For each of **(a)** the spin glass (SG) phase, $\delta = 1.5$; **(b)** the critical phase, $\delta = 0.0$; and **(c)**, the paramagnetic (PM) phase, $\delta = -1.5$, we generate a disorder realisation for $L = 20$ sites and apply RSRG-X. Additionally $g = 0.2$ and we use open boundary conditions. The columns respectively show: **(i)** The tree tensor network corresponding to each state. Bond decimations are shown as green rectangles linking two effective spins into one; site decimations are red circles which freeze an effective spin. The y-axis shows the step in which each decimation occurred, corresponding also to a rough energy scale. **(ii)** The z -components of the l-bits $\{\tau_k^z\}$, which can also be seen as the stabilisers of this state. Each row, corresponding to a decimation on the left, shows a single l-bit, with the colour giving the Pauli operator acting on each site. **(iii)** The x -components of the l-bits $\{\tau_k^x\}$, which can also be seen as the destabilisers.

full set of possible choices as building up a branching tree of approximate eigenstates [31]. (This is not to be confused with the tree tensor network structure of the states themselves.) For this reason, we refer to a full set of choices and the corresponding approximate eigenstate as a “leaf” of the RSRG-X tree. The geometry of the TTN depends on the set of choices made for decimation directions $\{c_k\}$. However, if we choose to fix the geometry, we can re-interpret these TTNs by considering the choice of decimation direction $c_k = \pm 1$ as an additional outgoing leg. In this picture,

decimation is represented by a long operator $-\sigma_l^y \sigma_{l+1}^z \dots \sigma_{r-1}^z \sigma_r^y$ across all the sites in the tree. The final site decimation in the SG phase typically has a very small energy scale associated with it, and may correspond to a strong zero mode linking the two symmetry sectors (see Sec. 4.3.1). Finally, the critical phase in the middle row (b) is a mixture of these two phases, containing both PM and SG regions. There is some freedom in how we define the l-bits, and in particular we may multiply any l-bit τ_k^z by some other $\tau_{k'}^z$ (which it must commute with) to obtain other valid l-bits.

The Clifford circuit representation allows us to efficiently calculate the action of any Pauli string (an operator that is the product of single-site Pauli operators) on a TTN state, by transforming it from the spin basis to the l-bit basis [234, 235]. Since the Clifford group maps Pauli strings to Pauli strings, this means such an operator maps one TTN state to exactly one other (with the same geometry), which may in fact be the same state. The Hamiltonian (4.1) is simply a sum of $O(L)$ Pauli strings which means that it retains this form in the l-bit basis, and so maps one TTN to at most $O(L)$ others. Therefore, we can efficiently calculate all matrix elements from a particular state, as well as expectation values.

The geometry of these tree tensor networks is largely informed by the balance of local fields and bonds, quantified by the value of δ . In order to better capture the effect of the interaction strength g , we also include the SW transformations in our wavefunction analysis. This is captured through an interaction picture: at each decimation the appropriate first-order SW transformation is calculated, $U_{\text{SW}} = \exp(iS^{(1)})$. We then apply the Clifford circuit, followed by SW transformations (expressed in the appropriate basis) up to first order to the Hamiltonian (or indeed to any operator), as $H^{(1)} = R^\dagger H R + [iS^{(1)}, R^\dagger H R]$. These SW transformations are analogous to the disentangling unitaries found in the multiscale entanglement renormalisation ansatz (MERA) [237], which sit in between layers of isometries on a TTN. By connecting adjacent branches, these capture short-ranged entanglement locally and enable MERA to efficiently describe entanglement at all length scales. Applying the Clifford circuit and SW transformations gives an effective Hamiltonian on the basis of l-bit product states, with those l-bits captured to first order. Each generator $S^{(1)}$ has support over a bounded number of sites, and so the effective Hamiltonian still has $O(L)$ terms. Despite this, computational complexity is still

significantly increased, limiting the maximum system size accessible to the order of hundreds of spins rather than thousands for the “zeroth-order” calculations. For full details of the Clifford RSRG-X method, including the application of the SW transformations, see Appendix 4.5.2.

The core output of our method therefore is a series of decimations, and associated SW transformations, which combine to create the localised state at a particular energy and also give us the effective local integrals of motion at that energy. These in turn are derived directly from the renormalised Hamiltonian at each step. Looking at the RG rules (Sec. 4.5.1), one can see that the terms in this Hamiltonian generally have contributions from multiple sources, and so at a late stage in the process the coefficients encode the detailed history of decimation choices. One can view these rules as summing up different processes that interfere with each other, constructively or destructively, and as a result the dominant term at some RG step is conditioned on previous choices. This also means that the geometry of the TTNs is energy-dependent.

This can be compared to the related spectrum bifurcation renormalisation group (SBRG) technique [106]. Unlike RSRG-X, SBRG aims to diagonalise the entire spectrum at once, extracting l-bit degrees of freedom but keeping these active (rotating them to σ^z operators) instead of removing them from the Hilbert space. This means crucially that the dominant term at any step only has contributions from a single process, without interference, and the RG flow is not energy-dependent. Our method, in comparison, is therefore able to capture more detailed information about the states, at the cost of requiring us to select a particular target energy.

4.3 Results

4.3.1 Spin-glass order

In Fig. 4.2(a) and (b), we show the scaling of the final l-bit’s energy ΔE_f (that is, the energy shift associated with the final RSRG-X decimation) with the parameter δ , averaged across disorder realisations and decimation choices, for two different values of g and a range of values of L . This should be the smallest energy scale in the system. Since the spin glass phase features spectral pairing [32], we should expect this energy difference to be exponentially small in system size L in the spin glass phase, but $O(1)$ in the paramagnetic phase. To test this, we plot $L^{-1/2} \ln |\Delta E_f|$ against $\delta L^{1/2}$ in panels (c) and (d). We observe a data collapse for both $g = 0$ and




Figure 4.2: (a, b) Energy scale of the final l-bit ΔE_f produced during the RSRG-X procedure, as a function of δ , for $g = 0$ and $g = 0.2$ respectively, and for various system sizes. (c, d) The same, but rescaled to show a clear data collapse and a phase transition at $\delta=0$ for all interaction strengths. The data is truncated (at the hollow circles) where $\ln|\Delta E_f| < -250$, as this is beyond the limits of numerical precision. Note since RSRG-X assumes strong randomness, we cannot find evidence of an ergodic phase directly from calculations like this.

$g = 0.2$, although the collapse is much cleaner for the non-interacting case, with the line tending quickly to zero for $\delta < 0$ and to a straight line through the origin for positive $\delta > 0$. (It is possible therefore that this collapse is not universal, but that the exponents of L here depend on the value of g .) Multiplying through by the factor of $L^{1/2}$, this clearly shows us that $\ln|\Delta E_f| \propto -L$ for fixed values of g and $\delta > 0$, agreeing with predictions. The final l-bit in the SG phase, to leading order, also always takes the form of two σ^y operators acting on either end of the largest spin cluster, with a string of σ^z operators in between. Expressed in terms Majorana fermions, this is in fact a bilocalised operator acting on the two fermions at either ends of this string. Deep into the SG phase, the largest cluster spans the system, so this becomes an edge mode, reminiscent of those found in superconducting quantum wires [226]. This operator anticommutes with the global parity operator $\prod_j \sigma_j^z$ which, given the exponentially small energy scale, this means that the final l-bit in the system in the SG phase is the strong zero mode (SZM) [128, 186, 189]. However, in the presence of interactions or in the marginal regime, sub-leading corrections become significant. We could use the Schrieffer-Wolff transformation to find these, but leave this to future work.

In applying RSRG-X we make an assumption of flow towards strong disorder,

and the infinite-randomness fixed point – implying the system is localised. Where we attempt to apply the method to a system which is in fact thermal, RSRG-X will produce inaccurate results. Previous work using exact, albeit small system size, numerical techniques [32, 33, 152], suggests that such a thermal phase exists for $\delta \simeq 0$. Hence the collapse found above, which would ordinarily tell us that the transition at $\delta=0$ becomes sharp in the thermodynamic limit, cannot be relied on as-is. What it *does* imply is that, where we do in fact have localisation and $\delta > 0$, we have an edge SZM. Therefore spectral pairing is always associated with the MBL SG phase, at least for the interaction strengths that we have looked at here.

If we carefully probe the results from RSRG-X for accuracy, we can instead show by contradiction that the assumption of flow to strong disorder is violated, and therefore that we are in an ergodic phase. In particular, we can use Clifford RSRG-X to investigate the approximate eigenstates produced as leaves of the RSRG-X tree, and their stability to off-diagonal parts of the Hamiltonian.

4.3.2 Energy variance

For each disorder realisation, we generate an RSRG-X leaf state $|\psi_0\rangle$ by picking a random set of decimation choices $\{c_k\}$ (that is, at infinite temperature), and then calculate the effective Hamiltonian on this state in the l-bit basis up to first order in SW transformations. We call this state the root state. This effective Hamiltonian is a sum of Pauli strings, with each Pauli string mapping the root state to exactly one state of definite l-bit configuration, $|\psi_\alpha\rangle$. We can then test the accuracy of the approximation by considering the “energy variance” [106],

$$\langle \delta H^2 \rangle = \langle \psi_0 | H^2 | \psi_0 \rangle - |\langle \psi_0 | H | \psi_0 \rangle|^2, \quad (4.6)$$

which measures to what extent the root state $|\psi_0\rangle$ is a good approximation of an actual eigenstate of the Hamiltonian H . It can also be written as $\sum_{\alpha \neq 0} |\langle \psi_0 | H | \psi_\alpha \rangle|^2$, leading to an interpretation of $\langle \delta H^2 \rangle$ as the degree to which the Hamiltonian maps the state away from itself. For perfect eigenstates, $\langle \delta H^2 \rangle = 0$.

Since $\langle \delta H^2 \rangle$ scales with the square of the total energy, which grows with system size and is not consistent for different choices of parameters, we choose to normalise this quantity by the square of the Hamiltonian *averaged across all states*,




Figure 4.3: Normalised energy variance $\langle \delta H^2 \rangle / \langle H^2 \rangle$ of RSRG-X leaf states across the δ - g plane, averaged over disorder realisations, for $L=500$. A perfect eigenstate has zero energy variance, but this quantity increases as the quality of approximation worsens. Since RSRG-X leaf states are localised, a large energy variance may indicate delocalisation and the approach to ergodicity. Also shown are selected contour lines (black).

$\overline{\langle H^2 \rangle} = \frac{1}{2L} \text{Tr } H^2$, which is basis independent. In Fig. 4.3, we plot this in the δ - g plane for a system of length $L=500$, showing that this quantity peaks towards the critical line $\delta=0$ and especially for larger values of g . One thing to note is that it appears to peak for small positive values of δ , rather than exactly at $\delta=0$, and similar results are seen for many of the other quantities we calculate – this is due to the open boundary conditions which leave fewer (spin-glass) order terms in the Hamiltonian relative to (paramagnetic) disorder terms.

The data show that the energy eigenstates in the critical region are less localised and cannot be well described by a single RSRG-X leaf state – it is likely that as we approach the critical line, the eigenstates pick up large fluctuations, and hence multiple l-bit basis states are needed to capture them. This may imply delocalisation of the eigenstates, but this depends on the details: if the number of basis states required grows extensively, then the system will thermalise in the thermodynamic limit, but otherwise these states will still occupy a vanishing fraction of the Hilbert space. To properly understand how the RSRG-X basis states hybridise to form the true eigenstates, we must therefore look at the resonances induced between them by off-diagonal Hamiltonian terms.

4.3.3 Many-body resonances

Each term of the transformed Hamiltonian maps the root state $|\psi_0\rangle$ to exactly one other tree tensor network state with the same geometry, which we label $|\psi_\alpha\rangle$. We can then calculate the many-body Thouless parameter [238] for each term,

$$\mathcal{G}_\alpha = \left| \frac{\langle \psi_0 | H | \psi_\alpha \rangle}{E_0 - E_\alpha} \right|. \quad (4.7)$$

When $\mathcal{G}_\alpha \ll 1$, the Hamiltonian only weakly couples the root state to nearby states in the Hilbert space, implying that the true eigenstate is close to the root state with only small contributions from other states at low orders in perturbation theory. However, as this quantity grows larger, perturbation theory begins to break down, with the root state becoming strongly resonant with other nearby states. This implies that the true eigenstates are superpositions of multiple states in the computational basis.

To make this a little more precise, for perturbation theory to converge the typical amplitude assigned to a diagram needs to decay faster than the combinatorial growth in the number of diagrams as the order of perturbation increases. We take $\mathcal{G}_\alpha$ as a rough estimate of this decay rate, and we choose to consider a resonance to have occurred when $\mathcal{G}_\alpha > \mathcal{G}^* = 0.1$. We believe this to be a slightly cautious threshold for what can be handled perturbatively. Any non-zero value of $\mathcal{G}_\alpha$ implies mixing of the tree tensor network basis; however, when small, we can take this to mean that the root state remains a good approximation of the true eigenstates up to some time, of order $1/\mathcal{G}^*$. Beyond that time perturbation theory would be required. In Sec. 4.5.4 we give some data for the alternative choice $\mathcal{G}^* = 1$, which is a much more optimistic threshold for what can be handled perturbatively.

The existence of a resonance does not necessarily mean that the eigenstates are no longer localised: when these resonances only connect a small number of states together (implying a small inverse participation ratio in the computational basis), the eigenstates may remain localised. Even if the number grows with system size, ergodicity is still avoided so long as the fraction of states involved vanishes in the thermodynamic limit. However, if resonances proliferate, they will connect an extensive number of states, leading to an ergodic phase. This is not the same as a thermal avalanche, wherein the resummation of one resonance creates a so-




Figure 4.4: Probability density of $\ln(\mathcal{G}_\alpha)$ for various points in the δ - g plane, on a logarithmic scale. The columns show (i) the PM phase with $\delta = -2$, (ii) the critical regime with $\delta = 0$, and (iii) the SG phase with $\delta = 2$, while the rows show (a) the non-interacting case $g = 0$, (b) $g = 0.05$, and (c) $g = 0.2$. The distributions peak at large values of $\mathcal{G}_\alpha$ when $\delta = 0$, with more strongly decaying left-hand tails indicating that the weight of said distributions are much more concentrated at these large values. Note that we only include states $|\psi_\alpha\rangle$ with nonzero matrix elements to $|\psi_0\rangle$. The resonance threshold $\mathcal{G} = 0.1$ is indicated with a dashed vertical line.

called superspin which, through the increased density of states, has an increased susceptibility to forming resonances with other l-bits. Here, we only consider the independent effects of off-diagonal terms and take clusters of directly resonant l-bits.

We show the distribution of $\mathcal{G}_\alpha$ across disorder realisations and root states in Fig. 4.4, for various values of δ and g . Close to the critical line $\delta = 0$ and with increasing interaction strength, these distributions peak more sharply (with faster-decaying left hand tails) at large values of $\mathcal{G}_\alpha$. Note that since each Hamiltonian term couples at most one other state to the root state, the total number of nonzero matrix elements is $O(L)$. Where two terms map to the same state, we combine their coefficients.

In Fig. 4.5(a), we count the number of terms for which the condition $\mathcal{G}_\alpha > \mathcal{G}^*$ is satisfied (that is, the number of resonances induced by the Hamiltonian), normalised




Figure 4.5: (a) Density (count per spin) of resonances with $\mathcal{G}_\alpha > \mathcal{G}^* = 0.1$, and (b) probability that a particular l-bit in a random disorder realisation is flipped by at least one such resonant term in the Hamiltonian. Both quantities calculated for randomly selected RSRG-X leaf states across the δ - g plane, averaged over disorder realisations, for $L=500$. Also shown are selected contours (black or white dashed lines). While the precise phase boundary between localisation and ergodicity cannot be located with this method, it is clear that the leaf states become unstable at increasingly small interaction strengths as one approaches the critical line $\delta=0$, in agreement with the findings of Ref. [32].

by the size of the system. Towards the critical line $\delta=0$ and at large (non-perturbative) interaction strengths g , we see this quantity peaking strongly, such that there is more than one resonance per spin in the system. This indicates a strong probability of crossover to an ergodic regime, although this does not allow us to locate the phase boundary precisely.

We also consider the chance that a particular l-bit will be flipped by at least one resonance in Fig. 4.5(b). This is subtly different to the average number of resonances per l-bit shown in (a) – in that it is less strongly influenced by rare regions with large numbers of resonant terms affecting a small subset of the l-bits. This peaks at a small positive value of $\delta=0$ and grows with increasing interaction strength g , to




Figure 4.6: Probability that a particular l-bit is in a resonant cluster of size s , for nine choices of parameters: $g = 0, 0.05, 0.2$ respectively in each row, and for models in the spin-glass, critical, and paramagnetic phases respectively in each column. This should tend to a constant in the thermodynamic limit. For this figure, a value of 0 indicates that an l-bit is unaffected by a resonant transition.

a greater than 90% chance – with almost every l-bit affected by a resonance, this makes it likely that the system is thermal in this portion of the phase diagram. To draw a more definitive conclusion however, we should look at the spatial distribution of these resonances.

4.3.4 Resonant Clusters

Consider the sets of l-bits respectively flipped by each resonant term. By taking the union of those sets with non-zero intersection, one arrives at a natural definition of “resonant clusters”: sets of degrees of freedom which are strongly mixed and locally thermal. This is analogous to percolation through a lattice where the links are formed by resonances. When resonances proliferate in a thermal phase [151, 152], these clusters grow to occupy a significant fraction of the system size.

In Fig. 4.6 we look at the distribution given by picking a random l-bit from a random disorder realisation and calculating the number of l-bits in (the size of) the




Figure 4.7: Maximum size of an l-bit cluster induced by resonances, averaged over disorder realisations, in the δ - g plane for $L=500$. (This quantity is considered to be 1 for a realisation with no resonances). The data here is shown on a logarithmic scale. Also shown are selected contour lines (white).

cluster it belongs to. (Here, an l-bit unaffected by resonances belongs to a cluster of size one.) When interactions are small and away from the critical line $\delta=0$, the l-bits form small clusters, and the probability of a larger cluster forming rapidly tails off, exponentially with cluster size. However, as interactions get stronger and we move towards $\delta=0$, the clusters grow larger, and in fact we can see these distributions are truncated by the finite system sizes accessible. For example, with $g = 0.2$ and $\delta=0$, we see the distribution peaking at the size of the system – it is overwhelmingly likely here that the resonances percolate through the entire system.

Fig. 4.7 then shows the per-realisation maximum size of these clusters, averaged over disorder, for a chain of length $L = 500$. Note that this is length-dependent since longer systems give more opportunities for large clusters to form. This appears to peak for small but positive δ , with the largest resonant clusters occupying almost the full system on average towards $g = 0.2$ and $\delta=0$. However, even at small interaction strengths, the largest cluster typically spans a substantial fraction of the system. In Sec. 4.5.3, we also show the probability distributions of the maximum cluster size over disorder realisations.

The data support the hypothesis that at the critical line the l-bits strongly hybridise and move the system towards an ergodic phase, even for smaller interaction strengths. To get a complete picture, we need to understand the scaling behaviour




Figure 4.8: (a) Disorder-averaged maximal resonant cluster size $\langle \max\{s_{\text{res}}\} \rangle$ against system size L , for $\delta=0$ (blue) and $\delta=-1$ (orange), each for $g=0$ (dots) and $g=0.2$ (triangles). Solid and dotted lines indicate power-law ($\langle \max\{s_{\text{res}}\} \propto L^\gamma$) or log-law ($\propto \log(L/L_0)$) fits respectively, to minimise the reduced- χ^2 statistic. (b) Fitted power-law exponent γ against δ for fixed $g=0.2$. The shaded region shows where the power-law fit is favoured over a log-law fit. Here, γ approaches 1, indicating the emergence of an extensive resonant cluster. (c) Difference in the reduced- χ^2 statistic between power-law and log-law fits, shown across the δ - g plane. For larger magnitude δ , the dependence on L grows weaker and the two hypotheses are more difficult to distinguish.

with system size. This will tell us whether resonances proliferate in the thermodynamic limit, indicating the breakdown of localisation as the dynamics become ergodic; or whether resonances grow much slower than the system size, such that each eigenstate only occupies a vanishing fraction of the Hilbert space.

4.3.5 Scaling of resonance behaviour with system size

In order to uncover the behaviour of the system in the thermodynamic limit, we need to understand how the resonant clusters scale with system size. To this end, in Fig. 4.8(a) we fit the mean maximal resonance size (see Figs. 4.6 and 4.7) to a power law in L , $\langle \max\{s_{\text{res}}\} \propto L^\gamma$, and to a log law, $\langle \max\{s_{\text{res}}\} \propto \log(L/L_0)$. In the localised phase, where resonances are rare and well-separated spatially, we should expect cluster sizes to be exponentially distributed and therefore $\langle \max\{s_{\text{res}}\} \rangle$ ought to scale as a logarithm, following the maximum order statistic of an exponential distribution. By contrast in the ergodic phase, when resonances proliferate, we should

expect the resultant clustering to percolate through the l-bits and form a cluster of size $O(L)$ such that $\gamma \rightarrow 1$. If the data follow a power law for smaller values of γ , this could still imply a power law in correlations and hence a diverging localisation length, even if the largest cluster does not extend across the entire system.

The data show that on the critical line $\delta=0$, the size of the maximum cluster scales as a power law with L , with $\gamma \simeq 1.0$. This implies that resonances proliferate such that the largest cluster occupies a finite fraction of the Hilbert space in the thermodynamic limit, destabilising the localised basis. Hence we conclude that the system is thermal at $\delta=0$. Looking away from the critical line, in Fig. 4.8(b), we show the fitted power-law exponent γ against δ , for fixed interaction strength $g = 0.2$. Additionally, we shade the region in which a power law is favoured over a log law. This shows that the power-law regime is accompanied by $\gamma \rightarrow 1$ and hence extensive resonant cluster scaling, verifying the intuition that this corresponds to a thermal phase.

Finally, Fig. 4.8(c) gives the difference between the reduced- χ^2 statistic between the power-law and log-law fits. The red region indicates that a power law is a better fit; the blue region corresponds to a log law. This shows clear evidence of an intervening ergodic phase, manifesting as a power law in resonant cluster scaling. The width of this phase grows with increasing g , but it is not entirely clear if this narrows to a single point for small but finite interaction strengths. (In our numerics, the smallest non-zero value we looked at was $g = 0.002$.)

4.4 Discussion

In this work we have developed a method to apply real-space renormalisation group to excited states (RSRG-X) of disordered spin-1/2 Hamiltonians and implicitly construct their wavefunctions as stabiliser states, even for large systems with many hundreds of spins, in the process also uncovering the l-bits for the system. This is done by constructing a Clifford circuit representing these approximate l-bits. Additionally, we have applied the Schrieffer-Wolff transformations to first order in order to improve accuracy, though at the cost of increased numerical complexity.

We have then applied this Clifford RSRG-X to the interacting Ising-Majorana chain, a model known to host two distinct MBL phases, and investigated the crossover between localised and ergodic behaviour in the supposed marginal MBL regime

between those two phases. By calculating the many-body Thouless parameter giving the strength of perturbative mixing between basis states, we show that resonances proliferate in the marginal MBL regime. This is shown to result in an intervening ergodic phase with boundaries similar to those found in studies of small systems [32–35].

Additionally, we have used Clifford RSRG-X to find the lowest-energy l-bit in the spin-glass phase. We show that this is a strong zero mode reminiscent of those found in superconducting quantum wires [226], with the leading term being a bilocalised Majorana fermion operator acting on either end of the largest spin cluster. Using this technique, it is possible to calculate higher-order corrections to this strong zero mode; however, we leave a systematic analysis to future work, instead focusing on the resonance picture and the ergodic regime in this work. Modifications of this technique may also allow access to higher-spin systems, enabling detailed characterisation of their MBL phases through determination of the l-bits and their higher-order corrections.

An ergodic phase has been argued to exist over an extended parameter regime even at arbitrarily weak interaction strength using ideas of a thermal avalanche, triggered by rare regions of weak disorder [32]. In contrast, we argue for an intervening thermal phase due a different mechanism unrelated to avalanches or rare regions. This is similar to the situation in small size numerics around the localisation transition, where MBL is destabilised despite the low likelihood of rare regions [239]. Still, avalanches may yet produce a wider ergodic phase than we find here. In the absence of a rigorous proof disallowing a direct MBL-MBL transition, it would be insightful to investigate other models of disorder such as stronger (power-law) disorder where this might occur. It is also noteworthy that resonance proliferation in this model seems to be confined to the region of the critical point, and directly heralds the breakdown of localisation.

Another promising direction would be to turn to quasiperiodic systems where rare regions do not occur, thus cannot precipitate an avalanche, and correlations in the disorder could be tunable independently of the transition. These effects could potentially stabilise MBL, as has been suggested for arrays of superconducting qubits [139], or even lead to a direct MBL-MBL transition. There has also been recent work

in constructing effective Hubbard models from continuous quasicrystalline models [140], and these may provide a more physically realistic testbed for these ideas than toy models such as the Aubry-Andre model.

Finally, we note that the Clifford RSRG-X machinery may be applied to other spin-1/2 models. For example, J.Z. Imbrie's proof of MBL [36] considers a transverse field Ising model with added longitudinal fields, limited to a regime where the transverse fields are perturbatively small; it would be interesting to consider variations on this model and see where localisation breaks down, especially given that the longitudinal fields break the $\mathbb{Z}_2$ symmetry protecting the Ising-Majorana model's topological phase. Chapter 5 describes exploratory work in that direction.

4.5 Supplementary Material

4.5.1 RSRG-X Decimation Rules

In these equations we consider any coupling that crosses the open boundary of the system to be zero. Additionally let J'_i and K_i be the strength of the term $\sigma_i^z \sigma_{i+1}^z$ and $\sigma_i^x \sigma_{i+2}^x$ respectively.

Site decimation rules: Suppose the largest gap is due to h_3 . Then we decimate site 3, setting $\sigma_3^z = c$, and renormalise the couplings as follows (with all unspecified couplings unaltered):

$$\tilde{h}_2 = h_2 + cJ'_2, \quad \tilde{h}_4 = h_4 + cJ'_3, \quad (4.8)$$

$$\begin{aligned} \tilde{J}_1 &= J_1 + c \frac{K_1 J_2}{h_3}, \quad \tilde{J}_2 = K_2 + c \frac{J_2 J_3}{h_3}, \\ \tilde{J}_4 &= J_4 + c \frac{K_3 J_3}{h_3}, \end{aligned} \quad (4.9)$$

$$\tilde{J}'_2 = 0, \quad \tilde{K}_1 = c \frac{K_1 J_3}{h_3}, \quad \tilde{K}_2 = c \frac{K_3 J_2}{h_3}. \quad (4.10)$$

We also calculate the change in the energy as,

$$\Delta E = c \left(h_3 + \frac{J_2^2 + J_3^2 + K_1^2 + K_3^2}{2h_3} \right). \quad (4.11)$$

Bond decimation rules: Suppose the largest gap is due to J_3 . Then we decimate the bond between sites 3 and 4, and the two sites are merged to create a new spin labelled c , renormalising the couplings as follows (with all unspecified couplings

unaltered):

$$\tilde{h}_2 = h_2 + c \frac{h_3 J'_2}{J_3}, \quad \tilde{h}_5 = h_5 + c \frac{h_4 J'_4}{J_3}, \quad (4.12)$$

$$\tilde{h}_c = J'_3 + c \frac{h_3 h_4}{J_3},$$

$$\tilde{J}_2 = c J_2 + K_2, \quad \tilde{J}_c = J_4 + c K_3, \quad (4.13)$$

$$\tilde{J}'_2 = c \frac{h_4 J'_2}{J_3}, \quad \tilde{J}'_c = c \frac{h_3 J'_4}{J_3}, \quad (4.14)$$

$$\tilde{K}_1 = c K_1, \quad \tilde{K}_2 = 0, \quad \tilde{K}_c = K_4. \quad (4.15)$$

We also calculate the change in the energy as,

$$\Delta E = c \left(J_3 + \frac{h_3^2 + h_4^2 + J_2'^2 + J_4'^2}{2h_3} \right). \quad (4.16)$$

4.5.2 Clifford RSRG-X

The starting point of the Clifford RSRG-X method is to apply traditional RSRG-X to a system, as per Ref. [31] – specifically, to a system described by a Hamiltonian in which each term is a Pauli string (a product of single-site Pauli operators). RSRG-X starts with the full system of L spins, and successively “decimates” degrees of freedom through the following prescription:

1. Locate the strongest term in the Hamiltonian $H_0 = \lambda A$, responsible for the largest energy gap.
2. Find and apply the Schrieffer-Wolff (SW) transformation e^{iS} which transforms the Hamiltonian to commute with H_0 – making the gap manifest.
3. Apply a Clifford transformation R to rotate H_0 to a Pauli σ_ℓ^z on some site ℓ , and decimate that site by freezing $\sigma_\ell^z = \pm 1$.
4. Return to step 1 and repeat until all degrees of freedom are frozen.

At each step, two transformations are generated: one, a Clifford rotation which maps Pauli strings to Pauli strings, and two, a Schrieffer-Wolff transformation which is more complicated. In order to analyse the properties of wavefunctions (and other related features, such as the effective Hamiltonian on the localised basis and matrix elements of operators), in many cases it has been sufficient to only include the Clifford rotations. One can combine these to form a Clifford circuit which prepares

the (approximate) localised basis from product states or, equivalently, maps operators on the physical spins to operators on the l-bits. Other RSRG-based methods have avoided additionally applying the SW transformation due to the increased complexity involved.

However, this process only obtains the localised basis to zeroth-order: each l-bit, which acts as a stabiliser to the basis, is a single Pauli string in the computational basis. In this work we found that this was not sufficient to capture the variation due to the interaction terms, motivating us to include the SW transformations in order to compute the localised basis to higher order. The transformation S at each step is given by the solution to,

$$[e^{iS}(H_0 + V)e^{-iS}, H_0] = 0 , \quad (4.17)$$

where $H = H_0 + V$. This can be expanded out and solved order-by-order in V . In the case of a Hamiltonian expressed in terms of Pauli strings, with a leading term $H_0 = \lambda A$ where A is a Pauli string, this is solved to first order by $S^{(0)} = 0$ and,

$$S^{(1)} = \frac{1}{4i\lambda^2} [H_0, V] . \quad (4.18)$$

Let us define e^{iS_i} and R_i to be the transformations at the i^{th} decimation step (here, we drop the superscript (1)). Then, we may write down the complete transformation which approximately diagonalises the Hamiltonian as,

$$\begin{aligned} U &= (R_L e^{iS_L})(R_{L-1} e^{iS_{L-1}}) \dots (R_2 e^{iS_2})(R_1 e^{S_1}) \\ &= (e^{i\tilde{S}_L} e^{i\tilde{S}_{L-1}} \dots e^{i\tilde{S}_1})(R_L R_{L-1} \dots R_1) \\ &\simeq e^{i\tilde{S}} R . \end{aligned} \quad (4.19)$$

Here, we have defined the notation,

$$\tilde{S}_i = R_{[i,L]} S_i R_{[i,L]}^\dagger , \quad (4.20)$$

$$R_{[i,L]} = R_L R_{L-1} \dots R_i , \quad (4.21)$$

such that $R_{[i,L]}$ is the partial Clifford circuit which transforms S_i (defined on the effective spins at that RG step) so that it instead acts upon the l-bit basis. In




Figure 4.9: Distribution of maximal resonant 1-bit cluster sizes for nine choices of parameters: $g = 0, 0.05, 0.2$ respectively in each row, for models in the spin-glass, critical, and paramagnetic phases respectively in each column. See Fig. 4.7 for comparison.

this way, we separate the transformation into two parts: a Clifford circuit R , and a product of unitaries $e^{i\tilde{S}}$. We are free here to treat the SW transformations as commuting, such that $e^A e^B = e^{A+B}$, since we are working to first order in V . We may then transform any operator (expressed as a sum of Pauli strings) to act upon the localised basis to first order by first pushing it through the Clifford circuit then applying the first-order SW transformation: $\hat{O} \rightarrow R^\dagger \hat{O} R + [i\tilde{S}, R^\dagger \hat{O} R]$. This is still expressed as a sum of Pauli strings, and so calculation of matrix elements etc. on 1-bit product states is straightforward – each Pauli string maps a product state to exactly one state (perhaps itself). Correspondingly, we can calculate the first-order 1-bits in the spin basis as $\tau^{x,z} = R[-i\tilde{S}, \sigma^{x,z}]R^\dagger$. For the purpose of implementing these calculations, we make use of the formalism in Ref. [234], representing Clifford circuits as binary matrices and Pauli strings as binary vectors (with an associated coefficient).

4.5.3 Distribution of maximal resonant cluster sizes

In Fig. 4.9 we show histograms across disorder realisations of the maximum resonant cluster size, $\max\{s_{\text{res}}\}$, for various system sizes L . The data show the maximum increasing with system size – this is to be expected, as a larger system gives more

chances for a large cluster to develop. Note also that in the thermal regime, the right-hand edge of the histograms is truncated by the system size with very little probability density on sizes smaller than this, showing that resonances dominate and are overwhelmingly likely to percolate throughout the entire system.

4.5.4 Higher resonance threshold

Throughout this work we have considered a resonance with $\mathcal{G} > \mathcal{G}^* = 0.1$ to be a resonance capable of destabilising the localised basis. In this section, we show some data for a higher threshold, $\mathcal{G}^* = 1$, meaning that we only consider very strong resonances which are certain to destabilise the basis. Fig. 4.10 is an analogue of Fig. 4.6, and shows the distribution of cluster sizes across disorder realisations. Additionally, Fig. 4.11 is an analogue of Fig. 4.9, and shows the distribution of *maximum* cluster sizes across disorder realisations.

Despite taking a much more conservative estimate of what is necessary to destabilise the system, there is still a trend towards delocalisation towards $\delta \simeq 0$ and $g \simeq 0.2$. The technique cannot (at present) be extended significantly beyond $g = 0.2$ because this would violate the assumption that the leading coupling comes from the relevant terms h_i and J_i , but the trend is clear and it seems likely that the scaling of the maximum cluster size would become extensive for larger g .




Figure 4.10: Similar to Fig. 4.6, but only considering very strong resonances with $\mathcal{G} > 1$.




Figure 4.11: Similar to Fig. 4.9, but only considering very strong resonances with $\mathcal{G} > 1$.

Chapter 5

Disordered Ising model with local fields

5.1 Introduction

While a rigorous proof of Anderson localisation in one dimension has existed for half a century [77], it was only a little under a decade ago that J.Z. Imbrie did the same for MBL in 1D [36]. This firmly established the phenomenology of MBL, with the existence of a complete set of quasi-local¹ integrals of motion [101, 102] which allow the system to avoid thermalisation indefinitely as detailed in Sec. 2.2.3. Specifically, the proof perturbatively constructs a unitary transformation which diagonalises the Hamiltonian. This unitary is itself a product of quasi-local rotations, and these are close to the identity except on certain sparse “resonant blocks” where perturbation theory breaks down.

While the formal proof of the existence of MBL was a remarkable achievement, there are a number of weaknesses in terms of the applicability. The assumption in the proof of limited energy-level attraction is reasonable given that all chaotic systems exhibit level repulsion while integrable systems (and numerically-studied MBL systems) show neutral Poisson statistics. However, the proof also strongly restricts the strength of the integrability breaking term in the model studied, in order to control resonances and maintain the validity of perturbation theory²; furthermore, the model has no symmetries besides time-reversal symmetry, and all terms are disordered. The proof should generalise to any model with similar properties, but this also precludes many of the models widely used in numerical and experimental studies of MBL, such as those with U(1) particle-conservation symmetry [67, 90, 96]

¹That is, with exponentially localised support.

²The exact expression in the proof is “ $\gamma^{1/20} \ll 1$ ”, where γ is the strength of the perturbation, implying γ is smaller than the 20th power of a small number!

and those with SPT order [31–35]. We are also usually interested in stronger integrability-breaking terms, particularly when considering the stability of MBL to external influences. For this last reason, in this chapter we choose to study the infinite-temperature phase diagram of the model studied in Ref. [36], for a broad range of parameters.

We introduce the precise model in the following Sec. 5.2, then in Sec. 5.3 we apply the Clifford RSRG-X machinery developed in the previous chapter. In Sec. 5.4, we apply perturbation theory to free-fermion states, with the intent of uncovering the extents of the MBL phase in the part of the phase diagram where the model is close to the transverse field Ising model (TFIM). Finally, in Sec. 5.5, we summarise and present a tentative phase diagram.

5.2 Model

The model studied in Ref. [36], up to a simple Clifford rotation, is given by:

$$H_{\text{Imbrie}} = \sum_{i=1}^{L-1} J_i \sigma_i^x \sigma_{i+1}^x + \sum_{i=1}^L h_i \sigma_i^x + \sum_{i=1}^L \Gamma_i \sigma_i^z. \quad (5.1)$$

We choose this definition as two of the three terms are the same as the dominant terms in H_{IM} (4.1), allowing us to re-use some of the machinery defined in the previous chapter. Every term in this model is disordered, and in Ref. [36] the transverse field Γ_i plays the role of a quantum perturbation, preventing the model from being classical and integrable. Γ_i is also deliberately chosen to be perturbatively small, in order to allow the proof to work.

While we are considering this model in the context of excited-state MBL, the ground state behaviour has been studied extensively, including using RSRG. Fig. 5.1 depicts the ground-state RG flow in H_{Imbrie} , with J_i , γ_i and h_i distributed i.i.d. in $[-1, 0]$, $[-\Gamma_0, 0]$, and $[-h_0/2, h_0/2]$ respectively, as found by Ref. [240]. The symmetric distribution for h ensures that there is no preferred longitudinal direction. With $h_0 = 0$, an infinite-randomness fixed point (IRFP) separates the quantum ordered and disordered phases of the TFIM; upon adding nonzero transverse fields, the quantum ordered phase is rendered unstable and now flows into the fixed point of the classical random-field Ising model, while the quantum disordered phase remains an attractive fixed point, with a separatrix extending from the IRFP. Of course, the




Figure 5.1: Schematic ground-state RG phase diagram of H_{Imbrie} , with J_i , Γ_i and h_i distributed i.i.d. in $[-1, 0]$, $[-\Gamma_0, 0]$, and $[-h_0/2, h_0/2]$ respectively. (Note that the ordering terms J_i are therefore ferromagnetic.) Arrows indicate the direction of RG flow. With $h_0 = 0$, an infinite-randomness fixed point (IRFP, red circle) separates the quantum ordered phase (black circle) from the quantum disordered phase (blue circle) of the TFIM. However, in the presence of longitudinal fields, the IRFP extends into a separatrix (red dashed line) between the quantum disordered phase and the fixed point of the classical random-field Ising model (green circle); the separatrix bends somewhat towards smaller Γ_0 with increasing h_0 . [From Ref. [240]. Reproduced under the [CC BY 4.0](#) Licence.]

ground state is not necessarily related to the excited-state phase diagram, and in particular provides no indication as to whether the spectrum will be many-body localised or not; on the other hand, MBL states resemble the ground states of gapped systems, so this may provide some indications as to what kind of states we can expect *if* the system remains localised. The model H_{Imbrie} can also be realised by placing the compound $\text{LiHo}_x\text{Y}_{1-x}\text{F}_4$ in a transverse magnetic field [241], and so has experimental relevance. Besides this exact model, random-field Ising models have long been studied in the context of phase transitions and infinite-randomness fixed points [242].

We expect from the analysis in Ref. [36] (see also Sec. 2.2.2) that when Γ_i is comparable to the other terms in the model, the system will be in an ergodic phase, but when small, it will be in an MBL phase. Likewise, when Γ_i is large, the model is a weakly perturbed version of H_{IM} from Eq. (4.1), and so the model should also be MBL. In fact, removing any one of the three terms makes the model integrable. With $J_i = 0$, the Hamiltonian is a sum of local fields, and therefore a trivial paramagnet with no interactions: a depth-1 unitary circuit (made up of local rotations at each site) will diagonalise this model. With $\Gamma_i = 0$, we have a classical Ising model, and the eigenstates are product states in the $\{\sigma_i^x\}$ basis. Depending on the relative strength of J_i and h_i , this will be a paramagnet or a ferromagnet. Finally, when $h_i = 0$, we

have the transverse field Ising model (TFIM), which is representable as free fermions via a Jordan-Wigner transform. Note also that with all three terms nonzero, there are *no symmetries* besides time reversal, and as such the only conserved quantity is energy. This means that there is only a single MBL phase, and no SPT order.

5.3 RSRG-X on Imbrie’s Model

The Hamiltonian H_{Imbrie} is very similar to the Ising-Majorana model H_{IM} of Eq. (4.1): both take the disordered transverse field Ising model as a root, but add different interacting terms. On one hand, the Ising-Majorana interaction terms are local four-body Majorana operators written in a spin language, and this affords the model a $\mathbb{Z}_2$ spin-flip symmetry (which corresponds to fermionic parity symmetry after a Jordan-Wigner transform). On the other hand, H_{Imbrie} uses longitudinal fields $\Gamma_i \sigma_i^x$ that explicitly break this symmetry, leaving only time reversal intact. This term is also both fermion-odd and non-local after a Jordan-Wigner transform, making interpretation in the fermion language tricky. However, as noted in the previous section, the model becomes integrable when any one of the three terms are removed, and so it is a matter of perspective which term is considered the “interaction” term. Due to the similarities between the models, an obvious first approach would be to reuse the RSRG-X machinery developed in the previous chapter (see Sec. 4.5.2), slightly modifying the decimation rules to deal with the altered interaction term. These derived decimation rules are laid out in the following section.

5.3.1 Decimation Rules

Bond decimation: The bond decimation, due to $H_0 = J_i \sigma_i^x \sigma_{i+1}^x$ is straightforward; in fact, because the longitudinal field commutes with H_0 here, it plays no direct role. The Schrieffer-Wolff transformation is given by:

$$S_{\text{Bond}}^{(1)} = - \left(\frac{\Gamma_i \sigma_i^y \sigma_{i+1}^x + \Gamma_{i+1} \sigma_i^x \sigma_{i+1}^y}{2J_i} \right). \quad (5.2)$$

We apply this to the effective Hamiltonian on the local four sites $\{i-1, i, i+1, i+2\}$ and then apply a Clifford rotation to bring $\sigma_i^x \sigma_{i+1}^x \rightarrow \sigma_i^z$, finally projecting out cross-terms and fixing $\sigma_i^z = c = \pm 1$. This gives the following effective Hamiltonian on

three sites $\{i-1, c, i+2\}$:

$$H_{\text{eff}} = cJ_i + \frac{c\Gamma_i^2}{2J_i} + \frac{c\Gamma_{i+1}^2}{2J_i} + h_{i-1}\sigma_{i-1}^x + (ch_i + h_{i+1})\sigma_c^x + h_{i+2}\sigma_{i+2}^x \\ + cJ_{i-1}\sigma_{i-1}^x\sigma_c^x + J_{i+1}\sigma_c^x\sigma_{i+2}^x + \frac{\Gamma_i\Gamma_{i+1}\sigma_c^z}{J_i} + \Gamma_{i-1}\sigma_{i-1}^z + \Gamma_{i+2}\sigma_{i+2}^z . \quad (5.3)$$

Therefore, the decimation rules are:

$$\Delta E = c \left(J_i + \frac{\Gamma_i^2 + \Gamma_{i+1}^2}{2J_i} \right) , \quad (5.4)$$

$$h_c \rightarrow ch_i + h_{i+1} , \quad J_{i-1} \rightarrow cJ_{i-1} , \quad J_c \rightarrow J_{i+1} , \quad \Gamma_c \rightarrow \frac{\Gamma_i\Gamma_{i+1}}{J_i} , \quad (5.5)$$

with all unspecified couplings unaltered (if still defined).

Site decimation: This is more complicated, because we have both a transverse field $\Gamma_i\sigma_i^z$ and a longitudinal field $h_i\sigma_i^x$. We could choose to decimate each type of field individually, here. However, it is possible to apply a single-site rotation, $U_{\text{rot}} = \exp(i\theta\sigma_i^y/2)$ with $\tan\theta = h_i/\Gamma_i$, to eliminate the longitudinal field, producing one single term $\Lambda_i\sigma_i^z$ with $\Lambda_i = \sqrt{h_i^2 + \Gamma_i^2}$. This is equivalent to treating these terms as a single off-axis field $\mathbf{\Lambda}_i \cdot \boldsymbol{\sigma}_i$ with $\mathbf{\Lambda}_i = (h_i, 0, \Gamma_i)^T$. Following this, we perform the same procedure as for the bond decimation. First, we find the Schrieffer-Wolff transformation,

$$S^{(1)} = \frac{\Gamma_i}{2\Lambda_i^2} (J_{i-1}\sigma_{i-1}^x\sigma_i^y + J_i\sigma_i^y\sigma_{i+1}^x) , \quad (5.6)$$

and then we apply this to the Hamiltonian on three local sites $\{i-1, i, i+1\}$, eliminating the central site i through a Clifford rotation $\sigma_i^x \rightarrow \sigma_i^z$ (specifically, the Hadamard gate) followed by projecting out cross-terms and fixing $\sigma_i^z = c = \pm 1$. This gives the following effective Hamiltonian on two sites $\{i-1, i+1\}$:

$$H_{\text{eff}} = c\Lambda_i + \frac{cJ_{i-1}^2\Gamma_i^2}{2\Lambda_i^3} + \frac{cJ_i^2\Gamma_i^2}{2\Lambda_i^3} + \left(h_{i-1} + \frac{ch_iJ_{i-1}}{\Lambda_i} \right) \sigma_{i-1}^x + \left(h_{i+1} + \frac{ch_iJ_i}{\Lambda_i} \right) \sigma_{i+1}^x \\ + \Gamma_{i-1}\sigma_{i-1}^z + \Gamma_{i+1}\sigma_{i+1}^z + \frac{cJ_{i-1}J_i\Gamma_i^2}{\Lambda_i^3} \sigma_{i-1}^x\sigma_{i+1}^x . \quad (5.7)$$




Figure 5.2: Probability distribution of (a) the maximal squared overlap between RSRG-X and ED, and (b) The inverse participation ratio of RSRG-X leaf states across the exact eigenstate basis. The legends in (a) and (b) also show, respectively, the arithmetic or geometric mean for each distribution. In both cases we sample 20 RSRG-X states in each of 100 disorder realisations, for $L = 10$.

Therefore the RG rules are:

$$\Delta E = c \left(\Lambda_i + \frac{\Gamma_i^2 (J_{i-1}^2 + J_i^2)}{2\Lambda_i^3} \right), \quad (5.8)$$

$$h_{i-1} \rightarrow h_{i-1} + \frac{ch_i J_{i-1}}{\Lambda_i}, \quad h_{i+1} \rightarrow h_{i+1} + \frac{ch_i J_i}{\Lambda_i}, \quad J_{i-1} \rightarrow \frac{c\Gamma_i^2 J_{i-1} J_i}{\Lambda_i^3}, \quad (5.9)$$

with all unspecified valid couplings unaltered.

5.3.2 Results

5.3.2.1 Validity of RSRG-X

We first must ensure that the RSRG-X procedure is valid for the model, has been implemented correctly, and determine the portions of the phase diagram over which it gives accurate results. The gold standard for quantum mechanical calculations is exact diagonalisation (ED), but it is only able to access relatively small system sizes. For this section, we pick i.i.d. coefficients from uniform distributions, with $J_i \in [0, J_0]$, $\Gamma_i \in [0, \Gamma_0]$, and $h_i \in [-h_0/2, h_0/2]$ – by picking $\langle h_i \rangle = 0$, we ensure that we break the $\mathbb{Z}_2$ symmetry of the TFIM in an unbiased manner.

For each of 30 disorder realisations on a chain of length $L = 10$, we generate 50 randomly-selected RSRG-X leaf states $|\psi_0\rangle$ and directly construct these, including applying the appropriate SW transformations at each step. We then compare these to the exact eigenstates $|n\rangle$, calculating the eigenstate decomposition of each leaf state $c_n = \langle n|\psi_0\rangle$. In Fig. 5.2, we show the histogram of (a) the maximal squared overlap for each leaf state, $\max_n |\langle n|\psi_0\rangle|^2$, and (b) the inverse participation ratio

(IPR) of each leaf state across the exact eigenstate basis, across a number of different parameter regimes. The IPR is defined as,

$$\text{IPR} = \sum_{n=1}^{\mathcal{D}} |\langle n|\psi\rangle|^4 = \overline{1/|c_n|^2}, \quad (5.10)$$

for a state $|\psi\rangle$ and some orthonormal basis $|n\rangle$ with dimension $\mathcal{D}$. The IPR approaches 1 for a state localised in the basis, and $1/\mathcal{D}$ for a state that is totally delocalised (or a random state) – therefore, we hope that it is close to 1 for leaf states in the eigenstate basis. Another interpretation is that $1/\text{IPR}$ is the number of basis states required to faithfully recreate the state. Together, these show that RSRG-X produces good approximations to the true eigenstates, with $\overline{1/\text{IPR}} \lesssim 2$ and $\overline{\max_n |c_n|^2} > 75\%$, so long as J_0 and Γ_0 are not large and comparable. That is to say, the procedure appears to work well in two of the three regimes described above: the trivial paramagnet where J_0 is small, and the classical Ising model where Γ_0 is small. However, the procedure appears to struggle in the TFIM limit by comparison (although it should be noted that this is a “limit” in a very loose sense: the small term here is quite large, $h_0 = 0.1$), and in the case where all three terms are large.

It is also possible to directly quantify the error produced by the RSRG-X algorithm. At each step, the Hamiltonian is rotated via a SW transformation into a basis which aligns the largest spectral gap with the largest term H_0 , followed by a projection into either the ground state or excited manifold of this term. At this point, any terms which anticommute with H_0 are discarded. We quantify the size of these terms by the operator norm of the commutator between the rotated Hamiltonian and $A = H_0/|H_0|$, up to second order in V ,

$$\delta_{H_{\text{eff}}} = \left\| \left[e^{iS^{(1)}} (H_0 + V) e^{-iS^{(1)}}, A \right] \right\|. \quad (5.11)$$

This must be calculated numerically, by finding exact solutions to the small system described by the non-commuting part. The procedure is valid as long as $\delta_{H_{\text{eff}}} \ll \Delta E$. Alternatively we may look at the sum of the squares of the coefficients in the Pauli string decomposition of this commutator, equivalent to the Frobenius norm, which




Figure 5.3: Probability distributions of the per-realisation maximum for (a) the operator errors $\delta_{H_{\text{eff}}}/H_0$, (b) the Frobenius errors $\varepsilon_{H_{\text{eff}}}/H_0$, and (c) the Schrieffer-Wolff coefficients $|S^{(1)}|$, for each possible choice of dominant terms. The distributions show the errors for the two cases (pink and orange) where both J_{XX} and Γ_Z are large and comparable. In all cases we choose $L = 1024$.




Figure 5.4: The disorder-averaged per-realisation maximum for (a) the operator errors $\delta_{H_{\text{eff}}}/H_0$, (b) the Frobenius errors $\varepsilon_{H_{\text{eff}}}/H_0$, and (c) the Schrieffer-Wolff coefficients $|S^{(1)}|$, against the size α of the small parameters in the Hamiltonian and for each possible choice of dominant terms. In all cases we choose $L = 1024$.

we call $\varepsilon_{H_{\text{eff}}}^2$. For a bond decimation, we have:

$$\varepsilon_{H_{\text{eff}}, \text{bond}}^2 = 4 \left(\frac{\Gamma_i^2 (h_i^2 + J_{i-1}^2) + \Gamma_{i+1}^2 (h_{i+1}^2 + J_{i+1}^2)}{J_i^2} \right), \quad (5.12)$$

and for a site decimation, we have:

$$\varepsilon_{H_{\text{eff}}, \text{site}}^2 = \frac{4\Gamma_i^2 [h_i^2 (J_{i-1}^2 + J_i^2)^2 + \Lambda_i^2 (J_{i-1}^2 \Gamma_{i-1}^2 + J_i^2 \Gamma_{i+1}^2)]}{\Lambda_i^6}. \quad (5.13)$$

This time taking $L = 1024$, we generate one RSRG-X leaf state per disorder

realisation, and compare the size of the off-diagonal terms produced at each step to the decimation energy scale H_0 , using $\delta_{H_{\text{eff}}}/|H_0|$ and $\varepsilon_{H_{\text{eff}}}/|H_0|$, and also look at the coefficients of the SW transform generators $|S^{(1)}|$, which we have implicitly assumed are small³. Therefore, in Fig. 5.3, we look at the distribution over disorder of the maximum value attained by each of these error measures in each single RSRG-X run, using the same parameter choices as in the comparison to ED. This confirms again that the renormalisation procedure is controlled, so long as at least one of J_0 or Γ_0 is small compared to the dominant term(s) in the Hamiltonian. In Fig. 5.4 we additionally take the geometric mean of these maximum values across disorder realisations, and show how these reduce when we change the size of the small terms. This shows that in the regime where both J_0 and Γ_0 are large, the RSRG-X errors do not decrease significantly when h_0 is made small, but in other regimes reducing the small terms likewise scales back the errors. This is despite the fact that the case with $h_0 = 0$ is *known* to be (Anderson) localised, and so it is reasonable to expect also that the $h_0 \ll 1$ case is MBL. This suggests that RSRG-X errors do not necessarily mean that the system delocalises in the regime, but it is clear we will need a different approach to successfully capture the eigenstates. We note also that the typical maximum error (not shown) is very close to the mean, implying that the mean is not strongly affected by rare disorder realisations (which could correspond to rare regions).

5.4 Fermion Perturbation Theory

When h_i is small, the model H_{Imbrie} is close to the TFIM, which can be solved exactly by a Jordan-Wigner transform to free fermions [243]. With J_i and Γ_i comparable, this is exactly the regime in which RSRG-X performs poorly, and suggests a better approach might be to start from the non-interacting TFIM limit and then treat the longitudinal fields $h_i\sigma_i^x$ as a perturbation.

³For a site decimation, this does *not* include the coefficient of the rotation U_{rot} , which is calculated and applied exactly

5.4.1 Method

We start by transforming the model H_{Imbrie} to free fermions, using a Jordan-Wigner transform:

$$\begin{aligned}\sigma_j^z &\longrightarrow a_j a_j^\dagger - a_j^\dagger a_j, \\ \sigma_j^x &\longrightarrow \left(\prod_{k \leq j} \sigma_k^z \right) (a_j^\dagger + a_j), \\ \sigma_j^x \sigma_{j+1}^x &\longrightarrow (a_j^\dagger - a_j) (a_{j+1}^\dagger + a_{j+1}),\end{aligned}\tag{5.14}$$

where we interpret $a_j^\dagger$ and a_j as the fermion creation and annihilation operators at site j . Hence, we can write the transformed Hamiltonian as $H_{\text{free}} + V$,

$$H_{\text{free}} = \sum_{j=1}^{j=L} \Gamma_j (a_j a_j^\dagger - a_j^\dagger a_j) + J_j (a_j^\dagger - a_j) (a_{j+1}^\dagger + a_{j+1}), \tag{5.15}$$

where V is our perturbation. If we further define Majorana fermion operators $r_{2j-1} = (a_j + a_j^\dagger)/\sqrt{2}$ and $r_{2j} = (a_j - a_j^\dagger)/\sqrt{2}i$, then we can write $H_{\text{free}} = i\mathbf{r}^\dagger h \mathbf{r}$ where h is a real antisymmetric matrix and $\mathbf{r}$ is the vector $(r_1, r_2, \dots, r_{2L})$. Crucially, this means we can diagonalise h via a real Schur decomposition, obtaining the free Majorana modes $s_j = \sum_k O_{jk} r_k$ given by the orthogonal matrix O_{jk} , and the energies $\{\varepsilon_j\}$ [243]. For a particular choice of occupied free modes, it is then easy to calculate the expectation value of any fermion-even observable using Wick's theorem. If ν_j is the occupation number of the j^{th} Dirac fermion mode, then we write down the Majorana covariance matrix in the free-fermion basis, γ_D , as

$$\gamma_D = \bigoplus_{j=1}^L \begin{pmatrix} 0 & \frac{1}{2} - \nu_j \\ \nu_j - \frac{1}{2} & 0 \end{pmatrix}, \tag{5.16}$$

and then transform this to the physical basis by $\gamma = O \gamma_D O^T$. Finally, the expectation value of some observable $V = r_{q_1} r_{q_2} \dots r_{q_N}$ is given by $\langle V \rangle = \text{Pf}(\gamma|_{\{q\}})$, where $\gamma|_{\{q\}}$ is the submatrix of γ on the sites $\{q\}$ and $\text{Pf}(A)$ is the Pfaffian.

The tricky part is then extending this to cover matrix elements, which are needed for perturbation theory calculations. Furthermore, $V = \sum_j h_j \sigma_j^x$ is fermion-odd and non-local (containing a long parity string to the left), which presents additional complications. Ref. [244] suggests adding a “ghost” (auxiliary) particle to the system,

which can be used to make the odd operators even, so long as the state of this extra particle is chosen carefully, while Ref. [245] is able to calculate entanglement entropies for a system with boundary fields (which are also fermion-odd). These do not work for our system, but there are simplifications we can make. We aim to calculate two quantities of interest using perturbation theory up to second order in V ,

$$\|\delta j\rangle\|^2 = \left\| \sum_{k \neq j} \frac{\langle k|V|j\rangle}{E_j - E_k} |k\rangle \right\|^2 = \sum_{k \neq j} \left| \frac{\langle k|V|j\rangle}{E_j - E_k} \right|^2, \quad (5.17)$$

$$\delta E = \langle j|V|j\rangle + \sum_{k \neq j} \frac{|\langle k|V|j\rangle|^2}{E_j - E_k}. \quad (5.18)$$

Because these only depend on the squared magnitude of the matrix element $|\langle k|V|j\rangle|^2$ and the expectation value $\langle j|V|j\rangle$, we do not need to worry about the phase of $|k\rangle$. We therefore can write $|k\rangle$ in terms of the initial state $|j\rangle$ as $|k\rangle = s_{\ell_1} s_{\ell_2} \dots s_{\ell_n} |j\rangle$, noting that an N -body V can flip at most N modes in $|j\rangle$, and so $n \leq N$ for any nonzero matrix elements. Letting $V = r_{q_1} r_{q_2} \dots r_{q_N}$ again, we can write it in the free-fermion basis as $V = \sum_{\{q'\}} O_{q_1 q'_1} O_{q_2 q'_2} \dots O_{q_N q'_N} s_{q'_1} s_{q'_2} \dots s_{q'_N}$, such that V is expressible as a rank- N tensor with dimension $2L$, $V_{q'_1 q'_2 \dots q'_N} = O_{q_1 q'_1} O_{q_2 q'_2} \dots O_{q_N q'_N}$. Hence,

$$\begin{aligned} \langle j|V|k\rangle &= \langle j|V s_{\ell_1} s_{\ell_2} \dots s_{\ell_n} |j\rangle \\ &= \sum_{\{q\}} \langle j|V_{q_1 q_2 \dots q_N} s_{q_1} s_{q_2} \dots s_{q_N} s_{\ell_1} s_{\ell_2} \dots s_{\ell_n} |j\rangle, \end{aligned} \quad (5.19)$$

which is an $\mathcal{O}((2L)^N)$ calculation when done exactly⁴. However, σ_L^x is $(2L-1)$ -body, which means this is super-exponential in system size if $V = \sum_j h_j \sigma_j^x$. For this reason, we decide to look at a slightly different perturbation,

$$V = \sum_{j=1}^{L-1} \lambda_j r_{2j-1} r_{2j} r_{2j+1} = \sum_{j=1}^{L-1} \lambda_j \left(\prod_{k < j} \sigma_k^z \right) \sigma_{j+1}^x. \quad (5.20)$$

This perturbation is fermion-odd and integrability breaking, just like $V = \sum_j h_j \sigma_j^x$, and so we hope that it will capture the same underlying physics as a longitudinal field.

⁴A smarter way of doing this may be to sample the matrix elements randomly, bounding the complexity but increasing uncertainty, but we leave this to future work.

Finally, we arrive at one minor complication. In the spin glass phase of the disordered TFIM, a strong zero mode imposes a spectrum-wide degeneracy, with the energy gap exponentially small in system size (see Sec. 4.3.1). As a result, we need to use nearly-degenerate perturbation theory: conceptually, one makes a minimal perturbation to the system to make the degeneracy exact, performs standard degenerate perturbation theory using V , and then reverses the initial perturbation in the new non-degenerate basis.

5.4.2 Results

Similarly to the model in Chapter 4, we take as our unperturbed system H_{free} (5.15), with Γ_j and J_j i.i.d. uniformly in $[0, \Gamma]$ and $[0, J]$ respectively. We then normalise by $\max(J, \Gamma) = 1$ and define the phase parameter $\delta = \ln J - \ln \Gamma$, such that the critical point is at $\delta = 0$ and $\delta > 0$ indicates the spin glass phase. Of course, V breaks the symmetry responsible for this phase transition, so we do not expect to see two different phases following the perturbation. When the state shift $\|\delta j\rangle\|$ becomes large, we expect that perturbation theory has failed to converge and the system has undergone a phase transition. For our purposes, we ignore any shift at 0th order due to V breaking the degeneracy, as this mixes precisely two states, and instead look only at the size of the first order shift. We consider two cases

- (a) $\lambda_1 = \lambda$ and all others zero (that is, the field acts only on the left edge, $V = \lambda_1 r_1 r_2 r_3$);
- (b) λ_j i.i.d. uniformly in $[-\lambda, \lambda]$.

In case (a), we calculate the critical perturbation strength λ^* by the conditions $\|\delta j\rangle\| = 1$, corresponding to there being at least one resonance induced by the perturbation. In case (b), we use the condition $\|\delta j\rangle\| = L$, which roughly corresponds to a resonance *density* of one. In both cases, we average over disorder realisations. We also considered an extended but uniform perturbation $\lambda_j = \lambda$, but this produced very similar results to case (b).

Fig. 5.5 depicts the dependence of λ^* on the phase parameter δ for various system sizes L and each of these two cases [respectively, Figs. 5.5(a) and 5.5(b)]. We can see that in both cases, we have a much smaller critical field near the critical point $\delta = 0$. There is also very different behaviour in the $\delta > 0$ regime between case (a), when the perturbation acts only on the left-hand edge of the system, and case




Figure 5.5: The disorder-averaged critical field λ^* (defined in the text), as a function of the tuning parameter δ and for a selected set of system sizes L , for each of two cases: **(a)** The perturbation V acts only on the left edge; **(b)** The perturbation is disordered and acts over the length of the chain. While there is uncertainty in λ^* , we suppress the error bars to avoid cluttering the figure.




Figure 5.6: The critical field λ^* against the inverse system size $1/L$, for selected values of δ , in each of the two cases (a) and (b). Error bars depict uncertainty due to disorder averaging. We also fit linear trendlines in $1/L$, allowing us to extrapolate the critical field as $L \rightarrow \infty$.

(b), where the perturbation is extended. In the former, the critical field grows very large for all L , indicating that the localised phase is stable to a perturbation of this kind, whereas in the latter case, the critical field stays constant. (We note that in actual fact, we would not be able to use perturbation theory when $\lambda \gg 1$, but we *can* conclude from $\lambda^* \gg 1$ that the system is stable to smaller perturbations.) However, we would also like to know the critical field in the thermodynamic limit, $\lim_{L \rightarrow \infty} \lambda^*$.

Therefore in Fig. 5.6, we show the finite-size scaling of λ^* for cases (a) and (b), plotting the critical field against $1/L$ for selected values of δ . We also fit the critical field to a straight line, allowing us to extrapolate to $L \rightarrow \infty$. Then, in Fig. 5.7, we




Figure 5.7: The extrapolated critical field λ^* in the thermodynamic limit $L \rightarrow \infty$, as a function of δ , for (a) a perturbation acting on the left edge and (b) an extended and disordered perturbation. This gives us a rough phase diagram, showing where we have either no resonances (green), some resonances (blue), and many ($\sim L$) resonances (orange).

plot this critical field as a function of δ , to obtain a phase diagram in each case. We see in Fig. 5.7(a) that an arbitrarily weak edge field can induce a resonance at the critical point $\Gamma = J$, but away from this point, the localised phase appears to remain intact for small field strengths. On the $\Gamma > J$ side, it appears that the critical field plateaus at $\lambda^* \simeq 0.7$, while when $J > \Gamma$, the critical field grows without bound. That being said, it is only for $\lambda \gtrsim 0.5$ that we see asymmetry between positive and negative δ , which is a relatively large value of λ where perturbation theory may not hold – and so the true phase diagram may look different to this. It is not clear whether this resonance would destabilise the MBL phase: it may cause a local thermal grain, which could potentially spread via a thermal avalanche, or it may just non-perturbatively change nearby l-bits.

On the other hand, in Fig. 5.7(b), we see that the perturbation induces sparse resonances for most portions of the phase diagram, except for a cone extending from the critical point $\delta = 0$, where we get a unit density of resonances (that is, $\gtrsim L$ resonances). While sparse resonances may or may not destabilise the MBL phase, this latter scenario is likely to cause a return to an ETH phase, as the resonances cause the state to be rearranged almost everywhere. Note that as well-separated parts of the chain do not communicate under the assumption of localisation, and each term in the perturbation has a finite chance of causing a resonance, there will always be $\mathcal{O}(L)$ resonances in the case of an extended perturbation, so there is no portion of the phase diagram with no resonances in the thermodynamic limit. Interestingly, the $\delta > 0$ side is more susceptible to resonances; this was also observed in Chapter 4.

5.5 Phase Diagram and Discussion

Finally, we combine our results to arrive at a provisional phase diagram from the model, in Fig. 5.8. When $h = 0$ and $\Gamma > J$, we have the paramagnetic phase of the TFIM, and this connects smoothly through the $\Gamma = J = 0$ point to the trivial paramagnet when $J = 0$ (where eigenstates are product states of spins aligned along the local fields). On the other hand, when $h = 0$ and $J > \Gamma$, we have the spin glass phase of the TFIM, which has SPT order (but of course this is broken immediately for $h > 0$). Finally, with $\Gamma = 0$, we have a classical Ising model, since the Hamiltonian only contains σ^x terms. Moving away from these integrable points, Fig. 5.4 tells us that the MBL phase is likely to be stable so long as one of J or Γ is small; that is, the trivial paramagnet remains localised for small J , while if we add small Γ to the classical Ising model, we also know from Ref. [36] that this will be MBL. We can therefore be fairly confident about these parts of the phase diagram.

From the results in Fig. 5.7(b), we see that the free-fermion localised states of the TFIM are unstable to an arbitrarily weak fully-disordered fermion-odd term at the critical point, and so we can conjecture that it will also be unstable to a perturbation $V = \sum_j h_j \sigma_j^x$. Near the critical point, the 3-fermion perturbation only induces some resonances, which may cause local ergodic grains to form. The MBL phase is expected to be unstable to thermal avalanches starting from such grains when the localisation length $\xi > 1/\ln 2$; such grains will also form in the thermodynamic limit due to rare low-disorder regions. Ref. [32] suggests that in the TFIM, the localisation length is $1/|\delta|$, which would mean the phase boundary intersects the TFIM at $\delta = \pm \ln 2$. Regardless, we would expect close proximity to the TFIM critical point to induce a large correlation length, and likely a large localisation length as well: this would ensure that an ETH phase would intervene with finite width sufficiently close to the transition. However, we have not determined the localisation length in our own calculations, and the extrapolation from the three-fermion term to a longitudinal field also introduces additional uncertainty.

We therefore leave the exact shape of the phase boundary to future research, particularly near the TFIM line. To answer this question, we could for example incorporate the resonance analysis from Chapter 4, or we could extend the perturbation theory technique with some other (to be determined) observable which better




Figure 5.8: The full schematic phase diagram for the disordered TFIM with longitudinal fields, shown on ternary axes. Solid, dashed, and dotted gridlines indicate lines of constant h , J and Γ respectively, and each corner represents the limit where one of these terms dominates. On the right-hand edge, with $\Gamma = 0$, we have a classical Ising model (green); on the left-hand edge, we have the trivial paramagnet with no nearest-neighbour term J (blue); and on the bottom edge, we have the non-interacting TFIM with $h = 0$. The TFIM is further split into the paramagnetic phase (blue) and the SPT-ordered spin glass phase (red), with a critical point (purple) between these. When one of Γ or J are small, we have an MBL phase (light blue); however, when Γ and J are both comparable and significant, we have an ETH phase (orange). This is expected to extend up to the TFIM line, but the exact intersection of the phase boundary with this line is unclear, indicated by dotted lines.

indicates the stability of the MBL phase. It is also possible that the states found by free-fermion diagonalisation could be used as a starting basis for an RSRG-X-like procedure, improving the accuracy: this would likely be able to handle non-local operators such as a longitudinal field in the fermion basis. We have also not determined the limits of the MBL phase in the regimes well-described by RSRG-X, and this would be a priority for obtaining a more accurate phase diagram. Further research would ideally also obtain the localisation length (or otherwise) in order to ascertain the susceptibility to thermal avalanches. However, it seems likely that the MBL phase in this model extends well beyond the very narrow limits set out in Imbrie’s proof of MBL.

Chapter 6

Conclusions

In this thesis, we have explored the interactions of both strong and weak ergodicity breaking with symmetry-protected topological order: how these novel out-of-equilibrium phenomena can stabilise quantum order, but also how quantum phase transitions can return a system to ergodicity.

In Chapter 3, we look at a certain spin-1/2 chain with $\mathbb{Z}_2 \times \mathbb{Z}_2$ SPT order that can be interpreted as a staggered ladder, which exhibits a long-lived coherent edge qubit when the even and odd sublattices are detuned [28]. In this same regime, we uncover towers of quantum many-body scars – atypical, low-entropy, ETH-violating states in an otherwise thermal spectrum – which preserve signatures of SPT order even at finite energy density, thus extending it out of the ground state. Furthermore, we show that quenches from initial states in the scarred manifold lead to bulk coherence in the dynamics. While previous studies were able to construct Hamiltonians which contained topologically-ordered scars in the middle of the spectrum [179, 180], these were finely-tuned models, whereas this result holds across a broad parameter range and should be stable to perturbations that do not break the symmetry. Furthermore, an approximation to this chain can be realised by periodically driving a 1D optical-tweezer array [28], raising the possibility of experimentally verifying the presence of these SPT-ordered scars.

In Chapter 4, we probe the phase transition between two MBL phases with different SPT order, at infinite temperature, in a spin-1/2 disordered Ising model with next-nearest neighbour interactions. This model can be mapped exactly onto a Majorana fermion chain with all possible two- and four-body local terms. Recent studies, albeit limited to small system sizes, found evidence in this model of an

intervening ergodic phase for arbitrarily weak interaction strengths, thus preventing a direct MBL-MBL transition [32–35]. At the same time, a series of recent results show many-body resonances and thermal avalanches [37–45] destroying MBL above a critical localisation length, and suggest that very high or even infinite disorder strength may be needed to stabilise MBL in the thermodynamic limit. We therefore studied the aforementioned SPT-MBL transition using a real space renormalisation group (RSRG-X) up to very large system sizes, showing that close to the transition many-body resonances overwhelm the localised phase and mix extensively many degrees of freedom. Crucially, we did not rely on any avalanche arguments to do so, showing that while thermal avalanches are important to understanding the phenomenology of MBL, they need not necessarily be involved in the breakdown of localisation. We also showed that one of the l-bits becomes the SPT strong zero edge mode in the topologically ordered phase. Our results have important implications for the nature of quantum order in the presence of disorder, suggesting that localisation (and the information stored in the l-bits) cannot survive when moving between different topological phases.

Finally, in Chapter 5, we study a similar model to the one in which J.Z. Imbrie proved the existence of MBL [36] using RSRG-X and methods from perturbation theory, with the aim of going beyond the narrow regime considered in that work. The ground state of this model has also previously been studied in the context of $\text{LiHo}_x\text{Y}_{1-x}\text{F}_4$, which gains an induced longitudinal field when placed in a transverse magnetic field [240, 241]. Our research uncovers an infinite-temperature phase diagram, showing that while MBL in this model appears stable in regimes where the non-interacting limit is a paramagnet or classical Ising model (and at much larger interaction strengths than considered in Ref. [36]), it is destabilised by proximity to the SPT critical point of the transverse field Ising model – even in the absence of the symmetry protecting this critical point. While this was only an exploratory study and further research is needed, it is intriguing to see how the proof of MBL holds up as the assumptions in it are pushed, and one pertinent question (in the context of recent interest in the breakdown of localisation) would be to explore exactly how the ergodic phase emerges in this simple model.

These studies, as well as the wider literature over the past few years, raise a

number of crucial questions. This thesis has focused on the situation in one dimension, where the only kind of topological order permitted is that protected by a symmetry (which also only results in short-range entanglement, rather than the long-range entanglement associated with true topological order). The use of ergodicity breaking to preserve topological order out of the ground state in two or more dimensions has been less well studied, and this would pose additional challenges. Unlike in 1D, where we have a rigorous proof of MBL given physically reasonable assumptions [36], in two or more dimensions it is not even clear if MBL exists at all as a proper phase of matter. True topological order would also require the protection of long-range entanglement, which would appear to be incompatible with localisation. However, quantum scars have been constructed in two and three dimensions with topological order [179, 180], and it would be interesting to see if topologically ordered scars could also be found in systems without fine tuning, similar to the SPT-ordered scars in Chapter 3. Furthermore, in four dimensions of space it is possible for topological order to extend to finite energy densities, even in clean systems without scarring [124, 246]. While we of course live in a three dimensional universe, certain quasicrystals (which may be realised e.g. in optical lattices [141]) may be interpreted as the projection of a higher-dimensional lattice and therefore could realise such exotic topological order.

One also wonders about alternative mechanisms to protect SPT order, beyond those explored in this thesis. Time-dependent driving appears to provide a route to novel out-of-equilibrium SPT order, including phases of matter that cannot be realised in thermal equilibrium – for example, those protected by discrete time translational symmetry [247], or which spontaneously break *multiple* time translational symmetries arising from a quasiperiodic drive [248]. Additionally, Ref. [249] shows that, even when SPT order is lost in static properties of the ground state, proximity to an SPT phase can lead to a measurable effect on the dynamics, and long-lived edge modes even within the symmetry broken phase. It is possible that some combination of these various phenomena with ergodicity breaking may lead to even further enhanced stability.

Topological order aside, the survival of MBL itself in the thermodynamic limit is a pressing question in the community. The discovery that rare low-disorder regions may precipitate thermal avalanches which then destabilise the entire MBL phase

[39] has thrown into doubt the previous consensus that MBL was a fairly generic phenomenon (somewhat like Anderson localisation), and the optimism following the celebrated proof of MBL in 1D [36]; now, some works question whether the disorder strength required to stabilise MBL is unreasonably large, or even infinite [40, 42, 43, 138]. Of course, no technical errors have been found in the proof, but it only claims validity in the regime of perturbatively small interaction strengths, and considers just one particular model with no symmetries, while numerical studies have considered a number of different models (e.g. with U(1) particle number conservation, or SPT order). And despite some promising experimental evidence [97], theory has somewhat dampened expectations of a stable MBL phase in $d > 1$ [250] – most importantly, it seems that MBL would always be unstable to thermal avalanches in higher dimensions, because the level spacing of an ergodic grain would fall *superexponentially* with size [39], and Ref. [135] suggests the proof of MBL would fail in $d > 1$. Since thermal avalanches are believed to spread through the mechanism of many-body resonances, the machinery developed in Chapter. 4 may be useful in answering some of these questions.

So is MBL and ergodicity breaking still relevant? Our experience of everyday life and classical physics suggests that avoiding thermalisation indefinitely should be impossible – the second law of thermodynamics always wins. The best we can hope for, outside of the idealised settings of perfectly isolated toy models (as per Ref. [36]), is therefore to protect systems from thermalisation over long timescales, and the evidence certainly suggests this is possible. Thermal avalanches, while destroying MBL in the infinite-time limit, should proceed on an extremely slow timescale [37, 134], thus allowing MBL to operate as a finite-time prethermal regime. Previous experimental results that appear to show MBL would not be invalidated by the recent arguments against the phenomenon: these only show the danger of extrapolating data for finite sizes and finite times out to the thermodynamic limit. We could also look to consider quasiperiodic disorder: this was not covered by Imbrie’s proof in Ref. [36], which only considered random disorder, but has the advantage of avoiding rare regions at the cost of inducing long-range correlations in disorder. Numerical and experimental evidence has shown that quasiperiodic disorder can induce MBL, and this may be a way to circumvent the avalanche phenomenon

[98, 139, 140, 251]. It would be simple to apply the Clifford RSRG-X method to a simple quasiperiodic model such as the Aubry-Andr  model, while methods similar to those developed in Chapter 5 could be used for e.g. the Hubbard models (with $U(1)$ particle conservation) in Ref. [140]. In the absence of a rigorous proof disallowing direct MBL-MBL transitions, it would also be interesting to see if quasiperiodic or power-law disorder may prevent an ergodic phase from intervening: these questions would also be amenable to Clifford RSRG-X. Meanwhile, recent discoveries such as quantum many-body scars prove that we still have a lot to learn about ergodicity breaking.

This research also has clear practical applications: today’s quantum technologies struggle with short decoherence timescales and shielding from noise, and qubits in a quantum computer do not need to survive forever – only long enough to fall below the error threshold for quantum error correcting codes to kick in. Lessons learnt from MBL and quantum scars should help overcome these barriers and achieve this goal: indeed, the use of disorder to stabilise qubits is an active area of research [252], and scars have been used to generate entangled states on quantum simulators [253]. In the other direction, the ergodicity-breaking community has a lot to learn from quantum information theory; in recent years, the state of the art in quantum many-body physics has incorporated a number of more advanced measures of entanglement, including mutual information, topological entanglement, and entanglement negativity [254], as well as new ways of thinking about thermalisation [255], and these all provide novel insights. Dual-unitary circuits also provide an exactly-solvable toy model of ergodicity and prethermalisation [256], which could lead to new ways of understanding these phenomena.

We should also consider some of the differences between the two ergodicity-breaking phenomena considered in this thesis, in terms of the state of the respective fields of study. Ever since Anderson first discovered single-particle localisation over half a century ago there have been predictions of localisation in a many-body setting, and since the first studies providing strong evidence of its existence [67, 90, 92–96] almost two decades ago, it has been an incredibly active area of research. There is also a general mechanism for producing MBL (namely, applying quenched disorder), allowing researchers to investigate a wide variety of systems with different properties.

This means that we now have a number of strong results about MBL, including a good understanding of the internal structure (through l-bits) and a proof of its existence. We are also now beginning to understand the limits of MBL, in particular whether it survives in the thermodynamic limit, and the nature of the prethermal localised regime in the case that it does not. On the other hand, quantum many-body scars were first detected in an experiment in 2017, making them a very new phenomenon. While there are specific mechanisms for inducing quantum scarring in systems (e.g. the projector-embedding form [17]), there is a strong element of fine tuning in many of these, and there are few generalisable results about scars. In particular, it is still an open question if there exist scars which are robust to perturbations in the thermodynamic limit. There are some frameworks which connect certain classes of scarred systems, such as Krylov-restricted thermalisation [156–158] or spectrum-generating algebras [15, 163], but it is still not clear if there is one single unifying structure to scars. Hopefully we will gain a clearer picture over the next few years: QMBS have attracted an incredible amount of interest as a form of *weak* ergodicity breaking, and there is still a lot to uncover. Likewise, regardless of its survival as a true phase of matter in the thermodynamic limit, MBL still holds relevance as a long-lived prethermal regime: rumours of its death have been greatly exaggerated.

Bibliography

- [1] J. Jeyaretnam, J. Richter, and A. Pal, “Quantum scars and bulk coherence in a symmetry-protected topological phase”, *Phys. Rev. B* **104**, 014424 (2021).
- [2] J. Jeyaretnam, C. J. Turner, and A. Pal, “Renormalization view on resonance proliferation between many-body localized phases”, *Phys. Rev. B* **108**, 094205 (2023).
- [3] M. Srednicki, “Chaos and quantum thermalization”, *Phys. Rev. E* **50**, 888 (1994).
- [4] M. Rigol, V. Dunjko, and M. Olshanii, “Thermalization and its mechanism for generic isolated quantum systems”, *Nature* **452**, 854 (2008).
- [5] L. D’Alessio, Y. Kafri, A. Polkovnikov, and M. Rigol, “From quantum chaos and eigenstate thermalization to statistical mechanics and thermodynamics”, *Adv. Phys.* **65**, 239 (2016).
- [6] J. M. Deutsch, “Eigenstate thermalization hypothesis”, *Rep. Prog. Phys.* **81**, 082001 (2018).
- [7] R. Nandkishore and D. A. Huse, “Many-Body Localization and Thermalization in Quantum Statistical Mechanics”, *Annu. Rev. Condens. Matter Phys.* **6**, 15 (2015).
- [8] D. A. Abanin, E. Altman, I. Bloch, and M. Serbyn, “Colloquium: Many-body localization, thermalization, and entanglement”, *Rev. Mod. Phys.* **91**, 021001 (2019).
- [9] H. Bernien, S. Schwartz, A. Keesling, H. Levine, A. Omran, H. Pichler, S. Choi, A. S. Zibrov, M. Endres, M. Greiner, V. Vuletić, and M. D. Lukin, “Probing many-body dynamics on a 51-atom quantum simulator.”, *Nature* **551**, 579 (2017).
- [10] S. Moudgalya, S. Rachel, B. A. Bernevig, and N. Regnault, “Exact excited states of nonintegrable models”, *Phys. Rev. B* **98**, 235155 (2018); S. Moudgalya, N. Regnault, and B. A. Bernevig, “Entanglement of exact excited states of Affleck-Kennedy-Lieb-Tasaki models: Exact results, many-body scars, and violation of the strong eigenstate thermalization hypothesis”, *Phys. Rev. B* **98**, 235156 (2018).
- [11] T. Iadecola and M. Schecter, “Quantum many-body scar states with emergent kinetic constraints and finite-entanglement revivals”, *Phys. Rev. B* **101**, 24306 (2020).
- [12] S. Chattopadhyay, H. Pichler, M. D. Lukin, and W. W. Ho, “Quantum many-body scars from virtual entangled pairs”, *Phys. Rev. B* **101**, 174308 (2020).
- [13] M. Schecter and T. Iadecola, “Weak Ergodicity Breaking and Quantum Many-Body Scars in Spin-1 XY Magnets”, *Phys. Rev. Lett.* **123**, 147201 (2019).
- [14] D. K. Mark, C.-J. Lin, and O. I. Motrunich, “Unified structure for exact towers of scar states in the Affleck-Kennedy-Lieb-Tasaki and other models”, *Phys. Rev. B* **101**, 195131 (2020).

- [15] N. O’Dea, F. Burnell, A. Chandran, and V. Khemani, “From tunnels to towers: Quantum scars from Lie algebras and q -deformed Lie algebras”, *Phys. Rev. Research* **2**, 043305 (2020).
- [16] S. Moudgalya, E. O’Brien, B. A. Bernevig, P. Fendley, and N. Regnault, “Large classes of quantum scarred Hamiltonians from matrix product states”, *Phys. Rev. B* **102**, 085120 (2020).
- [17] N. Shiraishi and T. Mori, “Systematic Construction of Counterexamples to the Eigenstate Thermalization Hypothesis”, *Phys. Rev. Lett.* **119**, 30601 (2017).
- [18] C. J. Turner, A. A. Michailidis, D. A. Abanin, M. Serbyn, and Z. Papić, “Weak ergodicity breaking from quantum many-body scars”, *Nat. Phys.* **14**, 745 (2018).
- [19] C. J. Turner, A. A. Michailidis, D. A. Abanin, M. Serbyn, and Z. Papić, “Quantum scarred eigenstates in a Rydberg atom chain: Entanglement, breakdown of thermalization, and stability to perturbations”, *Phys. Rev. B* **94**, 155134 (2018).
- [20] M. Serbyn, D. A. Abanin, and Z. Papić, “Quantum many-body scars and weak breaking of ergodicity”, *Nature Physics* **17**, 675 (2021).
- [21] I. Affleck, T. Kennedy, E. H. Lieb, and H. Tasaki, “Rigorous results on valence-bond ground states in antiferromagnets”, *Phys. Rev. Lett.* **59**, 799 (1987).
- [22] T. Kennedy and H. Tasaki, “Hidden symmetry breaking and the Haldane phase in $S = 1$ quantum spin chains”, *Commun. Math. Phys.* **147**, 431 (1992).
- [23] H. Li and F. D. M. Haldane, “Entanglement Spectrum as a Generalization of Entanglement Entropy: Identification of Topological Order in Non-Abelian Fractional Quantum Hall Effect States”, *Phys. Rev. Lett.* **101**, 010504 (2008).
- [24] F. Pollmann, A. M. Turner, E. Berg, and M. Oshikawa, “Entanglement spectrum of a topological phase in one dimension”, *Phys. Rev. B* **81**, 064439 (2010).
- [25] F. Pollmann, E. Berg, A. M. Turner, and M. Oshikawa, “Symmetry protection of topological phases in one-dimensional quantum spin systems”, *Phys. Rev. B* **85**, 075125 (2012).
- [26] W. Li, A. Weichselbaum, and J. v. Delft, “Identifying symmetry-protected topological order by entanglement entropy”, *Phys. Rev. B* **88**, 245121 (2013).
- [27] T. Senthil, “Symmetry-Protected Topological Phases of Quantum Matter”, *Annu. Rev. Condens. Matter Phys.* **6**, 299 (2015).
- [28] J. Kemp, N. Y. Yao, and C. R. Laumann, “Symmetry-Enhanced Boundary Qubits at Infinite Temperature”, *Phys. Rev. Lett.* **125**, 200506 (2020).
- [29] D. A. Huse, R. Nandkishore, V. Oganesyan, A. Pal, and S. L. Sondhi, “Localization-protected quantum order”, *Phys. Rev. B* **88**, 014206 (2013).
- [30] Y. Bahri, R. Vosk, E. Altman, and A. Vishwanath, “Localization and topology protected quantum coherence at the edge of hot matter”, *Nat. Commun.* **6**, 7341 (2015).
- [31] D. Pekker, G. Refael, E. Altman, E. Demler, and V. Oganesyan, “Hilbert-Glass Transition: New Universality of Temperature-Tuned Many-Body Dynamical Quantum Criticality”, *Phys. Rev. X* **4**, 011052 (2014).
- [32] N. Laflorencie, G. Lemarié, and N. Macé, “Topological order in random interacting Ising-Majorana chains stabilized by many-body localization”, *Phys. Rev. Res.* **4**, L032016 (2022).

- [33] S. Moudgalya, D. A. Huse, and V. Khemani, “Perturbative instability towards delocalization at phase transitions between MBL phases”, [arXiv:2008.09113 \[cond-mat.dis-nn\]](#) (2020).
- [34] R. Sahay, F. Machado, B. Ye, C. R. Laumann, and N. Y. Yao, “Emergent Ergodicity at the Transition between Many-Body Localized Phases”, [Phys. Rev. Lett. **126**, 100604 \(2021\)](#).
- [35] S. A. Parameswaran and R. Vasseur, “Many-body localization, symmetry and topology”, [Rep. Prog. Phys. **81**, 082501 \(2018\)](#).
- [36] J. Z. Imbrie, “On Many-Body Localization for Quantum Spin Chains”, [J. Stat. Phys. **163**, 998 \(2016\)](#); “Diagonalization and Many-Body Localization for a Disordered Quantum Spin Chain”, [Phys. Rev. Lett. **117**, 027201 \(2016\)](#).
- [37] K. Agarwal, E. Altman, E. Demler, S. Gopalakrishnan, D. A. Huse, and M. Knap, “Rare-region effects and dynamics near the many-body localization transition”, [Ann. Phys. \(Leipzig\) **529**, 1600326 \(2017\)](#).
- [38] D. J. Luitz, F. Huveneers, and W. De Roeck, “How a Small Quantum Bath Can Thermalize Long Localized Chains”, [Phys. Rev. Lett. **119**, 150602 \(2017\)](#).
- [39] W. De Roeck and F. Huveneers, “Stability and instability towards delocalization in many-body localization systems”, [Phys. Rev. B **95**, 155129 \(2017\)](#).
- [40] J. Šuntajs, J. Bonča, T. Prosen, and L. Vidmar, “Quantum chaos challenges many-body localization”, [Phys. Rev. E **102**, 062144 \(2020\)](#).
- [41] P. J. D. Crowley and A. Chandran, “Mean-field theory of failed thermalizing avalanches”, [Phys. Rev. B **106**, 184208 \(2022\)](#).
- [42] D. Sels, “Bath-induced delocalization in interacting disordered spin chains”, [Phys. Rev. B **106**, L020202 \(2022\)](#).
- [43] A. Morningstar, L. Colmenarez, V. Khemani, D. J. Luitz, and D. A. Huse, “Avalanches and many-body resonances in many-body localized systems”, [Phys. Rev. B **105**, 174205 \(2022\)](#).
- [44] D. M. Long, P. J. D. Crowley, V. Khemani, and A. Chandran, “Phenomenology of the Prethermal Many-Body Localized Regime”, [arXiv:2207.05761 \[cond-mat.dis-nn\]](#) (2022).
- [45] H. Ha, A. Morningstar, and D. A. Huse, “Many-Body Resonances in the Avalanche Instability of Many-Body Localization”, [Phys. Rev. Lett. **130**, 250405 \(2023\)](#).
- [46] T. Kinoshita, T. Wenger, and D. S. Weiss, “A quantum Newton’s cradle”, [Nature **440**, 900 \(2006\)](#).
- [47] U. Schneider, L. Hackermüller, J. P. Ronzheimer, S. Will, S. Braun, T. Best, I. Bloch, E. Demler, S. Mandt, D. Rasch, and A. Rosch, “Fermionic transport and out-of-equilibrium dynamics in a homogeneous Hubbard model with ultracold atoms”, [Nat. Phys. **8**, 213 \(2012\)](#).
- [48] M. V. Berry, “Regular and irregular semiclassical wavefunctions”, [J. Phys. A: Math. Gen. **10**, 2083 \(1977\)](#).
- [49] J. M. Deutsch, “Quantum statistical mechanics in a closed system”, [Phys. Rev. A **43**, 2046 \(1991\)](#).

- [50] J. R. Garrison and T. Grover, “Does a Single Eigenstate Encode the Full Hamiltonian?”, *Phys. Rev. X* **8**, 21026 (2018).
- [51] M. Rigol, “Breakdown of Thermalization in Finite One-Dimensional Systems”, *Phys. Rev. Lett.* **103**, 100403 (2009).
- [52] M. Rigol, “Quantum quenches and thermalization in one-dimensional fermionic systems”, *Phys. Rev. A* **80**, 053607 (2009).
- [53] R. Steinigeweg, J. Herbrych, and P. Prelovšek, “Eigenstate thermalization within isolated spin-chain systems”, *Phys. Rev. E* **87**, 012118 (2013).
- [54] W. Beugeling, R. Moessner, and M. Haque, “Finite-size scaling of eigenstate thermalization”, *Phys. Rev. E* **89**, 042112 (2014).
- [55] H. Kim, T. N. Ikeda, and D. A. Huse, “Testing whether all eigenstates obey the Eigenstate Thermalization Hypothesis”, *Phys. Rev. E* **90**, 52105 (2014).
- [56] R. Mondaini, K. R. Fratus, M. Srednicki, and M. Rigol, “Eigenstate thermalization in the two-dimensional transverse field Ising model”, *Phys. Rev. E* **93**, 032104 (2016).
- [57] D. Jansen, J. Stolpp, L. Vidmar, and F. Heidrich-Meisner, “Eigenstate thermalization and quantum chaos in the Holstein polaron model”, *Phys. Rev. B* **99**, 155130 (2019).
- [58] M. Brenes, T. LeBlond, J. Goold, and M. Rigol, “Eigenstate Thermalization in a Locally Perturbed Integrable System”, *Phys. Rev. Lett.* **125**, 070605 (2020).
- [59] J. Richter, A. Dymarsky, R. Steinigeweg, and J. Gemmer, “Eigenstate thermalization hypothesis beyond standard indicators: Emergence of random-matrix behavior at small frequencies”, *Phys. Rev. E* **102**, 042127 (2020); J. Wang, M. H. Lamann, J. Richter, R. Steinigeweg, A. Dymarsky, and J. Gemmer, “Eigenstate Thermalization Hypothesis and Its Deviations from Random-Matrix Theory beyond the Thermalization Time”, *Phys. Rev. Lett.* **128**, 180601 (2022); A. Dymarsky, “Bound on Eigenstate Thermalization from Transport”, *Phys. Rev. Lett.* **128**, 190601 (2022).
- [60] R. Kubo, “Statistical-Mechanical Theory of Irreversible Processes. I. General Theory and Simple Applications to Magnetic and Conduction Problems”, *J. Phys. Soc. Jpn.* **12**, 570 (1957).
- [61] P. C. Martin and J. Schwinger, “Theory of Many-Particle Systems. I”, *Phys. Rev.* **115**, 1342 (1959).
- [62] E. Khatami, G. Pupillo, M. Srednicki, and M. Rigol, “Fluctuation-Dissipation Theorem in an Isolated System of Quantum Dipolar Bosons after a Quench”, *Phys. Rev. Lett.* **111**, 050403 (2013).
- [63] L. Foini and J. Kurchan, “Eigenstate thermalization hypothesis and out of time order correlators”, *Phys. Rev. E* **99**, 042139 (2019).
- [64] S. Pappalardi, L. Foini, and J. Kurchan, “Eigenstate Thermalization Hypothesis and Free Probability”, *Phys. Rev. Lett.* **129**, 170603 (2022).
- [65] O. Bohigas, M. J. Giannoni, and C. Schmit, “Characterization of Chaotic Quantum Spectra and Universality of Level Fluctuation Laws”, *Phys. Rev. Lett.* **52**, 1 (1984).
- [66] F. Haake, Random-Matrix Theory, in *Quantum Signatures of Chaos* (Springer Berlin Heidelberg, Berlin, Heidelberg, 2001) pp. 47–118.
- [67] V. Oganesyan and D. A. Huse, “Localization of interacting fermions at high temperature”, *Phys. Rev. B* **75**, 155111 (2007).

- [68] Y. Y. Atas, E. Bogomolny, O. Giraud, and G. Roux, “Distribution of the Ratio of Consecutive Level Spacings in Random Matrix Ensembles”, *Phys. Rev. Lett.* **110**, 84101 (2013).
- [69] J. M. Deutsch, “Thermodynamic entropy of a many-body energy eigenstate”, *New J. Phys.* **12**, 075021 (2010); J. M. Deutsch, H. Li, and A. Sharma, “Microscopic origin of thermodynamic entropy in isolated systems”, *Phys. Rev. E* **87**, 042135 (2013).
- [70] A. M. Kaufman, M. E. Tai, A. Lukin, M. Rispoli, R. Schittko, P. M. Preiss, and M. Greiner, “Quantum thermalization through entanglement in an isolated many-body system”, *Science* **353**, 794 (2016).
- [71] P. Calabrese and J. Cardy, “Evolution of entanglement entropy in one-dimensional systems”, *J. Stat. Mech.: Theory Exp.* **2005** (04), P04010.
- [72] P. Calabrese and J. Cardy, “Entanglement entropy and quantum field theory”, *J. Stat. Mech.: Theory Exp.* **2004** (06), P06002.
- [73] H. Castella, X. Zotos, and P. Prelovšek, “Integrability and Ideal Conductance at Finite Temperatures”, *Phys. Rev. Lett.* **74**, 972 (1995).
- [74] J. Sirker, R. G. Pereira, and I. Affleck, “Conservation laws, integrability, and transport in one-dimensional quantum systems”, *Phys. Rev. B* **83**, 035115 (2011).
- [75] P. W. Anderson, “Absence of Diffusion in Certain Random Lattices”, *Phys. Rev.* **109**, 1492 (1958).
- [76] R. Abou-Chacra, D. J. Thouless, and P. W. Anderson, “A selfconsistent theory of localization”, *J. Phys. C* **6**, 1734 (1973).
- [77] V. L. Berezinskiĭ, Kinetics of a quantum particle in a one-dimensional random potential, in *30 Years of the Landau Institute — Selected Papers* (World Scientific Series in 20th Century Physics, 1996) pp. 87–94.
- [78] E. Abrahams, P. W. Anderson, D. C. Licciardello, and T. V. Ramakrishnan, “Scaling Theory of Localization: Absence of Quantum Diffusion in Two Dimensions”, *Phys. Rev. Lett.* **42**, 673 (1979).
- [79] D. S. Wiersma, P. Bartolini, A. Lagendijk, and R. Righini, “Localization of light in a disordered medium”, *Nature* **390**, 671 (1997).
- [80] J. Billy, V. Josse, Z. Zuo, A. Bernard, B. Hambrecht, P. Lugan, D. Clément, L. Sanchez-Palencia, P. Bouyer, and A. Aspect, “Direct observation of Anderson localization of matter waves in a controlled disorder”, *Nature* **453**, 891 (2008).
- [81] J. Chabé, G. Lemarié, B. Grémaud, D. Delande, P. Szriftgiser, and J. C. Garreau, “Experimental observation of the anderson metal-insulator transition with atomic matter waves”, *Phys. Rev. Lett.* **101**, 255702 (2008).
- [82] S. Karbasi, C. R. Mirr, P. G. Yarandi, R. J. Frazier, K. W. Koch, and A. Mafi, “Observation of transverse Anderson localization in an optical fiber”, *Opt. Lett.* **37**, 2304 (2012).
- [83] S. Karbasi, R. J. Frazier, K. W. Koch, T. Hawkins, J. Ballato, and A. Mafi, “Image transport through a disordered optical fibre mediated by transverse Anderson localization”, *Nat. Commun.* **5**, 3362 (2014).
- [84] T. Ying, Y. Gu, X. Chen, X. Wang, S. Jin, L. Zhao, W. Zhang, and X. Chen, “Anderson localization of electrons in single crystals: $\text{Li}_x\text{Fe}_7\text{Se}_8$ ”, *Sci. Adv.* **2**, e1501283 (2016).

- [85] A. Yamilov, S. E. Skipetrov, T. W. Hughes, M. Minkov, Z. Yu, and H. Cao, “Anderson localization of electromagnetic waves in three dimensions”, *Nat. Phys.*, **1** (2023).
- [86] K. Chadán, N. N. Khuri, A. Martin, and T. Tsun Wu, “Bound states in one and two spatial dimensions”, *J. Math. Phys.* **44**, 406 (2003).
- [87] D. J. Thouless, “Localization distance and mean free path in one-dimensional disordered systems”, *J. Phys. C* **6**, L49 (1973).
- [88] H. Hu, A. Strybulevych, J. H. Page, S. E. Skipetrov, and B. A. van Tiggelen, “Localization of ultrasound in a three-dimensional elastic network”, *Nat. Phys.* **4**, 945 (2008).
- [89] N. F. Mott, “Conduction in non-crystalline materials”, *Philos. Mag. (1798-1977)* **19**, 835 (1969).
- [90] A. Pal and D. A. Huse, “The many-body localization phase transition”, *Phys. Rev. B* **82**, 174411 (2010).
- [91] B. L. Altshuler, Y. Gefen, A. Kamenev, and L. S. Levitov, “Quasiparticle Lifetime in a Finite System: A Nonperturbative Approach”, *Phys. Rev. Lett.* **78**, 2803 (1997).
- [92] I. V. Gornyi, A. D. Mirlin, and D. G. Polyakov, “Interacting electrons in disordered wires: Anderson localization and low-T transport”, *Phys. Rev. Lett.* **95**, 206603 (2005).
- [93] D. M. Basko, I. L. Aleiner, and B. L. Altshuler, “Metal–insulator transition in a weakly interacting many-electron system with localized single-particle states”, *Ann. Phys. (NY)* **321**, 1126 (2006).
- [94] L. B. Ioffe and M. Mézard, “Disorder-Driven Quantum Phase Transitions in Superconductors and Magnets”, *Phys. Rev. Lett.* **105**, 037001 (2010).
- [95] C. Gogolin, M. P. Müller, and J. Eisert, “Absence of Thermalization in Nonintegrable Systems”, *Phys. Rev. Lett.* **106**, 040401 (2011).
- [96] M. Žnidarič, T. Prosen, and P. Prelovšek, “Many-body localization in the Heisenberg XXZ magnet in a random field”, *Phys. Rev. B* **77**, 064426 (2008).
- [97] J. Choi, S. Hild, J. Zeiher, P. Schauß, A. Rubio-Abadal, T. Yefsah, V. Khemani, D. A. Huse, I. Bloch, and C. Gross, “Exploring the many-body localization transition in two dimensions”, *Science* **352**, 1547 (2016).
- [98] M. Schreiber, S. S. Hodgman, P. Bordia, H. P. Lüschen, M. H. Fischer, R. Vosk, E. Altman, U. Schneider, and I. Bloch, “Observation of many-body localization of interacting fermions in a quasi-random optical lattice”, *Science* **349**, 842 (2015).
- [99] J. Smith, A. Lee, P. Richerme, B. Neyenhuis, P. W. Hess, P. Hauke, M. Heyl, D. A. Huse, and C. Monroe, “Many-body localization in a quantum simulator with programmable random disorder”, *Nat. Phys.* **12**, 907 (2016).
- [100] R. K. Panda, A. Scardicchio, M. Schulz, S. R. Taylor, and M. Žnidarič, “Can we study the many-body localisation transition?”, *Europhys. Lett.* **128**, 67003 (2020).
- [101] D. A. Huse, R. Nandkishore, and V. Oganesyan, “Phenomenology of fully many-body-localized systems”, *Phys. Rev. B* **90**, 174202 (2014).
- [102] J. Z. Imbrie, V. Ros, and A. Scardicchio, “Local integrals of motion in many-body localized systems”, *Ann. Phys. (Leipzig)* **529**, 1600278 (2017).

- [103] R. Vosk and E. Altman, “Many-Body Localization in One Dimension as a Dynamical Renormalization Group Fixed Point”, *Phys. Rev. Lett.* **110**, 067204 (2013).
- [104] R. Vasseur, A. C. Potter, and S. A. Parameswaran, “Quantum Criticality of Hot Random Spin Chains”, *Phys. Rev. Lett.* **114**, 217201 (2015).
- [105] A. Goremykina, R. Vasseur, and M. Serbyn, “Analytically Solvable Renormalization Group for the Many-Body Localization Transition”, *Phys. Rev. Lett.* **122**, 040601 (2019).
- [106] Y.-Z. You, X.-L. Qi, and C. Xu, “Entanglement holographic mapping of many-body localized system by spectrum bifurcation renormalization group”, *Phys. Rev. B* **93**, 104205 (2016).
- [107] D. Pekker, B. K. Clark, V. Oganesyan, and G. Refael, “Fixed Points of Wegner-Wilson Flows and Many-Body Localization”, *Phys. Rev. Lett.* **119**, 075701 (2017).
- [108] S. J. Thomson and M. Schiró, “Time evolution of many-body localized systems with the flow equation approach”, *Phys. Rev. B* **97**, 060201(R) (2018).
- [109] S. J. Thomson and M. Schirò, “Local integrals of motion in quasiperiodic many-body localized systems”, *SciPost Phys.* **14**, 125 (2023).
- [110] F. Pollmann, V. Khemani, J. I. Cirac, and S. L. Sondhi, “Efficient variational diagonalization of fully many-body localized Hamiltonians”, *Phys. Rev. B* **94**, 041116(R) (2016).
- [111] T. B. Wahl, A. Pal, and S. H. Simon, “Efficient Representation of Fully Many-Body Localized Systems Using Tensor Networks”, *Phys. Rev. X* **7**, 021018 (2017).
- [112] T. B. Wahl, F. Venn, and B. Béri, “Local integrals of motion detection of localization-protected topological order”, *Phys. Rev. B* **105**, 144205 (2022).
- [113] A. Chandran, J. Carrasquilla, I. H. Kim, D. A. Abanin, and G. Vidal, “Spectral tensor networks for many-body localization”, *Phys. Rev. B* **92**, 024201 (2015).
- [114] D. Pekker and B. K. Clark, “Encoding the structure of many-body localization with matrix product operators”, *Phys. Rev. B* **95**, 035116 (2017).
- [115] A. Chandran, I. H. Kim, G. Vidal, and D. A. Abanin, “Constructing local integrals of motion in the many-body localized phase”, *Phys. Rev. B* **91**, 085425 (2015).
- [116] M. Goihl, M. Gluza, C. Krumnow, and J. Eisert, “Construction of exact constants of motion and effective models for many-body localized systems”, *Phys. Rev. B* **97**, 134202 (2018).
- [117] A. K. Kulshreshtha, A. Pal, T. B. Wahl, and S. H. Simon, “Behavior of l-bits near the many-body localization transition”, *Phys. Rev. B* **98**, 184201 (2018).
- [118] V. Ros, M. Müller, and A. Scardicchio, “Integrals of motion in the many-body localized phase”, *Nuc. Phys. B* **891**, 420 (2015).
- [119] J. H. Bardarson, F. Pollmann, and J. E. Moore, “Unbounded growth of entanglement in models of many-body localization”, *Phys. Rev. Lett.* **109**, 17202 (2012).
- [120] M. Serbyn, Z. Papić, and D. A. Abanin, “Universal Slow Growth of Entanglement in Interacting Strongly Disordered Systems”, *Phys. Rev. Lett.* **110**, 260601 (2013).
- [121] M. Serbyn, Z. Papić, and D. A. Abanin, “Quantum quenches in the many-body localized phase”, *Phys. Rev. B* **90**, 174302 (2014).

- [122] R. Orús, “A practical introduction to tensor networks: Matrix product states and projected entangled pair states”, *Ann. Phys. (NY)* **349**, 117 (2014).
- [123] M. Serbyn, A. A. Michailidis, D. A. Abanin, and Z. Papić, “Power-Law Entanglement Spectrum in Many-Body Localized Phases”, *Phys. Rev. Lett.* **117**, 160601 (2016).
- [124] B. Bauer and C. Nayak, “Area laws in a many-body localized state and its implications for topological order”, *J. Stat. Mech.: Theory Exp.* **2013** (09), P09005.
- [125] R. Vasseur, S. A. Parameswaran, and J. E. Moore, “Quantum revivals and many-body localization”, *Phys. Rev. B* **91**, 140202 (2015).
- [126] M. Serbyn, M. Knap, S. Gopalakrishnan, Z. Papić, N. Y. Yao, C. R. Laumann, D. A. Abanin, M. D. Lukin, and E. A. Demler, “Interferometric Probes of Many-Body Localization”, *Phys. Rev. Lett.* **113**, 147204 (2014).
- [127] M. C. Bañuls, N. Y. Yao, S. Choi, M. D. Lukin, and J. I. Cirac, “Dynamics of quantum information in many-body localized systems”, *Phys. Rev. B* **96**, 174201 (2017).
- [128] P. Fendley, “Strong zero modes and eigenstate phase transitions in the XYZ/interacting Majorana chain”, *J. Phys. A* **49**, 30LT01 (2016).
- [129] R. Nandkishore, S. Gopalakrishnan, and D. A. Huse, “Spectral features of a many-body-localized system weakly coupled to a bath”, *Phys. Rev. B* **90**, 064203 (2014).
- [130] E. Altman, “Many-body localization and quantum thermalization”, *Nat. Phys.* **14**, 979 (2018).
- [131] W. De Roeck, F. Huveneers, M. Müller, and M. Schiulaz, “Absence of many-body mobility edges”, *Phys. Rev. B* **93**, 014203 (2016).
- [132] J. Léonard, S. Kim, M. Rispoli, A. Lukin, R. Schittko, J. Kwan, E. Demler, D. Sels, and M. Greiner, “Probing the onset of quantum avalanches in a many-body localized system”, *Nat. Phys.* **19**, 481 (2023).
- [133] R. B. Griffiths, “Nonanalytic Behavior Above the Critical Point in a Random Ising Ferromagnet”, *Phys. Rev. Lett.* **23**, 17 (1969).
- [134] T. Thiery, F. Huveneers, M. Müller, and W. De Roeck, “Many-Body Delocalization as a Quantum Avalanche”, *Phys. Rev. Lett.* **121**, 140601 (2018).
- [135] W. D. Roeck and J. Z. Imbrie, “Many-body localization: stability and instability”, *Phil. Trans. R. Soc. A* **375**, 20160422 (2017).
- [136] L. Colmenarez, D. J. Luitz, and W. De Roeck, “Ergodic inclusions in many body localized systems”, arXiv [10.48550/arXiv.2308.01350](https://arxiv.org/abs/10.48550/arXiv.2308.01350) (2023), [2308.01350](https://arxiv.org/abs/2308.01350).
- [137] J. C. Peacock and D. Sels, “Many-body delocalization from embedded thermal inclusion”, *Phys. Rev. B* **108**, L020201 (2023).
- [138] D. Sels and A. Polkovnikov, “Dynamical obstruction to localization in a disordered spin chain”, *Phys. Rev. E* **104**, 054105 (2021).
- [139] E. Varvelis and D. P. DiVincenzo, “Perturbative Analysis of Quasi-periodic Patterning of Transmon Quantum Computers: Enhancement of Many-Body Localization”, [arXiv:2212.03805 \[quant-ph\]](https://arxiv.org/abs/2212.03805) (2022).
- [140] E. Gottlob and U. Schneider, “Hubbard models for quasicrystalline potentials”, *Phys. Rev. B* **107**, 144202 (2023).

- [141] J.-C. Yu, S. Bhave, L. Reeve, B. Song, and U. Schneider, “Observing the two-dimensional Bose glass in an optical quasicrystal”, [arXiv:2303.00737 \[cond-mat.quant-gas\]](#) (2023).
- [142] D. Sels and A. Polkovnikov, “Thermalization of Dilute Impurities in One-Dimensional Spin Chains”, [Phys. Rev. X](#) **13**, 011041 (2023).
- [143] S. J. Garratt, S. Roy, and J. T. Chalker, “Local resonances and parametric level dynamics in the many-body localized phase”, [Phys. Rev. B](#) **104**, 184203 (2021).
- [144] V. Khemani, S. P. Lim, D. N. Sheng, and D. A. Huse, “Critical Properties of the Many-Body Localization Transition”, [Phys. Rev. X](#) **7**, 021013 (2017).
- [145] R. Nandkishore and A. C. Potter, “Marginal Anderson localization and many-body delocalization”, [Phys. Rev. B](#) **90**, 195115 (2014).
- [146] D. A. Abanin, J. H. Bardarson, G. De Tomasi, S. Gopalakrishnan, V. Khemani, S. A. Parameswaran, F. Pollmann, A. C. Potter, M. Serbyn, and R. Vasseur, “Distinguishing localization from chaos: Challenges in finite-size systems”, [Ann. Phys.](#) **427**, 168415 (2021).
- [147] E. J. Heller, “Bound-State Eigenfunctions of Classically Chaotic Hamiltonian Systems: Scars of Periodic Orbits”, [Phys. Rev. Lett.](#) **53**, 1515 (1984).
- [148] D. Wintgen and A. Hönig, “Irregular wave functions of a hydrogen atom in a uniform magnetic field”, [Phys. Rev. Lett.](#) **63**, 1467 (1989).
- [149] E. J. Heller, “Wavepacket dynamics and quantum chaology”, in [Chaos and Quantum Physics](#), Vol. 52, edited by M. Giannoni, A. Voros, and J. Z. Justin (1991).
- [150] P. B. Wilkinson, T. M. Fromhold, L. Eaves, F. W. Sheard, N. Miura, and T. Takamasu, “Observation of ‘scarred’ wavefunctions in a quantum well with chaotic electron dynamics”, [Nature](#) **380**, 608 (1996).
- [151] I. V. Protopopov, W. W. Ho, and D. A. Abanin, “Effect of SU(2) symmetry on many-body localization and thermalization”, [Phys. Rev. B](#) **96**, 041122(R) (2017).
- [152] I. V. Protopopov, R. K. Panda, T. Parolini, A. Scardicchio, E. Demler, and D. A. Abanin, “Non-Abelian Symmetries and Disorder: A Broad Nonergodic Regime and Anomalous Thermalization”, [Phys. Rev. X](#) **10**, 011025 (2020).
- [153] P. Zhang, Y.-C. Lai, *et al.*, “Many-body Hilbert space scarring on a superconducting processor”, [Nat. Phys.](#) **19**, 120 (2023).
- [154] G.-X. Su, H. Sun, A. Hudomal, J.-Y. Desaulles, Z.-Y. Zhou, B. Yang, J. C. Halimeh, Z.-S. Yuan, Z. Papić, and J.-W. Pan, “Observation of many-body scarring in a Bose-Hubbard quantum simulator”, [Phys. Rev. Res.](#) **5**, 023010 (2023).
- [155] Á. M. Alhambra, A. Anshu, and H. Wilming, “Revivals imply quantum many-body scars”, [Phys. Rev. B](#) **101**, 205107 (2020).
- [156] S. Moudgalya, A. Prem, R. Nandkishore, N. Regnault, and B. A. Bernevig, “Thermalization and Its Absence within Krylov Subspaces of a Constrained Hamiltonian”, in [Memorial Volume for Shoucheng Zhang](#) (World Scientific, Singapore, 2020) pp. 147–209.
- [157] P. Sala, T. Rakovszky, R. Verresen, M. Knap, and F. Pollmann, “Ergodicity Breaking Arising from Hilbert Space Fragmentation in Dipole-Conserving Hamiltonians”, [Phys. Rev. X](#) **10**, 011047 (2020).

- [158] K. Lee, A. Pal, and H. J. Changlani, “Frustration-induced emergent Hilbert space fragmentation”, *Phys. Rev. B* **103**, 235133 (2021).
- [159] S. Moudgalya, N. Regnault, and B. A. Bernevig, “ η -pairing in Hubbard models: From spectrum generating algebras to quantum many-body scars”, *Phys. Rev. B* **102**, 085140 (2020).
- [160] V. Khemani, C. R. Laumann, and A. Chandran, “Signatures of integrability in the dynamics of Rydberg-blockaded chains”, *Phys. Rev. B* **99**, 161101 (2019).
- [161] Z. Yao, L. Pan, S. Liu, and H. Zhai, “Quantum many-body scars and quantum criticality”, *Phys. Rev. B* **105**, 125123 (2022).
- [162] A. Daniel, A. Hallam, J.-Y. Desaulles, A. Hudomal, G.-X. Su, J. C. Halimeh, and Z. Papić, “Bridging quantum criticality via many-body scarring”, *Phys. Rev. B* **107**, 235108 (2023).
- [163] S. Choi, C. J. Turner, H. Pichler, W. W. Ho, A. A. Michailidis, Z. Papić, M. Serbyn, M. D. Lukin, and D. A. Abanin, “Emergent SU(2) Dynamics and Perfect Quantum Many-Body Scars”, *Phys. Rev. Lett.* **122**, 220603 (2019).
- [164] D. Yuan, S.-Y. Zhang, and D.-L. Deng, “Exact Quantum Many-Body Scars in Higher-Spin Kinetically Constrained Models”, *arXiv:2307.06357 [cond-mat.str-el]* (2023).
- [165] N. Shibata, N. Yoshioka, and H. Katsura, “Onsager’s Scars in Disordered Spin Chains”, *Phys. Rev. Lett.* **124**, 180604 (2020).
- [166] I. Mondragon-Shem, M. G. Vavilov, and I. Martin, “Fate of Quantum Many-Body Scars in the Presence of Disorder”, *PRX Quantum* **2**, 030349 (2021).
- [167] C. Chen, F. Burnell, and A. Chandran, “How Does a Locally Constrained Quantum System Localize?”, *Phys. Rev. Lett.* **121**, 085701 (2018).
- [168] P. Sierant, E. G. Lazo, M. Dalmonte, A. Scardicchio, and J. Zakrzewski, “Constraint-Induced Delocalization”, *Phys. Rev. Lett.* **127**, 126603 (2021).
- [169] S.-Y. Zhang, D. Yuan, T. Iadecola, S. Xu, and D.-L. Deng, “Extracting Quantum Many-Body Scarred Eigenstates with Matrix Product States”, *Phys. Rev. Lett.* **131**, 020402 (2023).
- [170] P. Kramer, “A review of the time-dependent variational principle”, *J. Phys. Conf. Ser.* **99**, 012009 (2008).
- [171] J. Haegeman, C. Lubich, I. Oseledets, B. Vandereycken, and F. Verstraete, “Unifying time evolution and optimization with matrix product states”, *Phys. Rev. B* **94**, 165116 (2016).
- [172] L. A. Bunimovich, “On the ergodic properties of nowhere dispersing billiards”, *Commun. Math. Phys.* **65**, 295 (1979).
- [173] W. W. Ho, S. Choi, H. Pichler, and M. D. Lukin, “Periodic Orbits, Entanglement, and Quantum Many-Body Scars in Constrained Models: Matrix Product State Approach”, *Phys. Rev. Lett.* **122**, 040603 (2019).
- [174] A. Hallam, J.-Y. Desaulles, and Z. Papić, “Embedding semiclassical periodic orbits into chaotic many-body Hamiltonians”, *arXiv:2303.01359 [quant-ph]* (2023).
- [175] C. J. Turner, J.-Y. Desaulles, K. Bull, and Z. Papić, “Correspondence Principle for Many-Body Scars in Ultracold Rydberg Atoms”, *Phys. Rev. X* **11**, 021021 (2021).

- [176] X. Chen, Z.-C. Gu, and X.-G. Wen, “Classification of gapped symmetric phases in one-dimensional spin systems”, *Phys. Rev. B* **83**, 035107 (2011).
- [177] A. Chandran, V. Khemani, C. R. Laumann, and S. L. Sondhi, “Many-body localization and symmetry-protected topological order”, *Phys. Rev. B* **89**, 144201 (2014).
- [178] J. A. Kjäll, J. H. Bardarson, and F. Pollmann, “Many-Body Localization in a Disordered Quantum Ising Chain”, *Phys. Rev. Lett.* **113**, 107204 (2014).
- [179] S. Ok, K. Choo, C. Mudry, C. Castelnovo, C. Chamon, and T. Neupert, “Topological many-body scar states in dimensions one, two, and three”, *Phys. Rev. Res.* **1**, 33144 (2019).
- [180] N. S. Srivatsa, J. Wildeboer, A. Seidel, and A. E. B. Nielsen, “Quantum many-body scars with chiral topological order in two dimensions and critical properties in one dimension”, *Phys. Rev. B* **102**, 235106 (2020).
- [181] X. Chen, Z.-C. Gu, and X.-G. Wen, “Complete classification of one-dimensional gapped quantum phases in interacting spin systems”, *Phys. Rev. B* **84**, 235128 (2011).
- [182] A. M. Turner, F. Pollmann, and E. Berg, “Topological phases of one-dimensional fermions: An entanglement point of view”, *Phys. Rev. B* **83**, 075102 (2011).
- [183] L. Fidkowski and A. Kitaev, “Effects of interactions on the topological classification of free fermion systems”, *Phys. Rev. B* **81**, 134509 (2010).
- [184] L. Fidkowski and A. Kitaev, “Topological phases of fermions in one dimension”, *Phys. Rev. B* **83**, 075103 (2011).
- [185] F. Pollmann and A. M. Turner, “Detection of symmetry-protected topological phases in one dimension”, *Phys. Rev. B* **86**, 125441 (2012).
- [186] P. Fendley, “Parafermionic edge zero modes in $\mathbb{Z}_n$ -invariant spin chains”, *J. Stat. Mech.: Theory Exp.* **2012** (11), P11020.
- [187] X. Mi, M. Sonner, M. Y. Niu, K. W. Lee, B. Foxen, H. Neven, D. Bacon, J. Hilton, E. Lucero, R. Babbush, S. Boixo, A. Megrant, Y. Chen, J. Kelly, V. Smelyanskiy, D. A. Abanin, P. Roushan, *et al.*, “Noise-resilient edge modes on a chain of superconducting qubits”, *Science* **378**, 785 (2022).
- [188] V. V. Sivak, A. Eickbusch, B. Royer, S. Singh, I. Tsioutsios, S. Ganjam, A. Miano, B. L. Brock, A. Z. Ding, L. Frunzio, S. M. Girvin, R. J. Schoelkopf, and M. H. Devoret, “Real-time quantum error correction beyond break-even”, *Nature* **616**, 50 (2023).
- [189] J. Kemp, N. Y. Yao, C. R. Laumann, and P. Fendley, “Long coherence times for edge spins”, *J. Stat. Mech.: Theory Exp.* **2017** (6), 063105.
- [190] D. Abanin, W. De Roeck, W. W. Ho, and F. Huveneers, “A Rigorous Theory of Many-Body Prethermalization for Periodically Driven and Closed Quantum Systems”, *Commun. Math. Phys.* **354**, 809 (2017).
- [191] D. V. Else, P. Fendley, J. Kemp, and C. Nayak, “Prethermal Strong Zero Modes and Topological Qubits”, *Phys. Rev. X* **7**, 041062 (2017).
- [192] A. C. Potter and R. Vasseur, “Symmetry constraints on many-body localization”, *Phys. Rev. B* **94**, 224206 (2016).
- [193] F. Venn, T. B. Wahl, and B. Béri, “Many-body-localization protection of eigenstate topological order in two dimensions”, *arXiv:2212.09775 [cond-mat.dis-nn]* (2022).

- [194] J. Wildeboer, A. Seidel, N. S. Srivatsa, A. E. B. Nielsen, and O. Erten, “Topological quantum many-body scars in quantum dimer models on the kagome lattice”, *Phys. Rev. B* **104**, L121103 (2021).
- [195] W. Son, L. Amico, R. Fazio, A. Hamma, S. Pascazio, and V. Vedral, “Quantum phase transition between cluster and antiferromagnetic states”, *EPL (Europhysics Lett.)* **95**, 50001 (2011).
- [196] H. J. Briegel and R. Raussendorf, “Persistent Entanglement in Arrays of Interacting Particles”, *Phys. Rev. Lett.* **86**, 910 (2001).
- [197] A. C. Doherty and S. D. Bartlett, “Identifying Phases of Quantum Many-Body Systems That Are Universal for Quantum Computation”, *Phys. Rev. Lett.* **103**, 020506 (2009).
- [198] D. V. Else, S. D. Bartlett, and A. C. Doherty, “Symmetry protection of measurement-based quantum computation in ground states”, *New J. Phys.* **14**, 113016 (2012).
- [199] J. K. Pachos and M. B. Plenio, “Three-Spin Interactions in Optical Lattices and Criticality in Cluster Hamiltonians”, *Phys. Rev. Lett.* **93**, 056402 (2004).
- [200] S. D. Bartlett and T. Rudolph, “A simple nearest-neighbor two-body Hamiltonian system for which the ground state is a universal resource for quantum computation”, *Phys. Rev. A* **74**, 040302 (2006).
- [201] T. Griffin and S. D. Bartlett, “Spin lattices with two-body Hamiltonians for which the ground state encodes a cluster state”, *Phys. Rev. A* **78**, 062306 (2008).
- [202] R. Verresen, R. Moessner, and F. Pollmann, “One-dimensional symmetry protected topological phases and their transitions”, *Phys. Rev. B* **96**, 165124 (2017).
- [203] P. Pfeuty, “The one-dimensional Ising model with a transverse field”, *Ann. Phys. (NY)* **57**, 79 (1970).
- [204] A. W. Sandvik, “Computational Studies of Quantum Spin Systems”, *AIP Conf. Proc.* **1297**, 135 (2010).
- [205] P. Weinberg and M. Bukov, “QuSpin: a Python Package for Dynamics and Exact Diagonalisation of Quantum Many Body Systems part I: spin chains”, *SciPost Phys.* **2**, 003 (2017).
- [206] P. Weinberg and M. Bukov, *QuSpin 0.3.4* (2020).
- [207] D. N. Page, “Average entropy of a subsystem”, *Phys. Rev. Lett.* **71**, 1291 (1993).
- [208] J. Eisert, M. Cramer, and M. B. Plenio, “Colloquium: Area laws for the entanglement entropy”, *Rev. Mod. Phys.* **82**, 277 (2010).
- [209] A. Lavasani, Y. Alavirad, and M. Barkeshli, “Measurement-induced topological entanglement transitions in symmetric random quantum circuits”, *Nat. Phys.* , 342 (2021).
- [210] M. V. Berry and M. Tabor, “Level clustering in the regular spectrum”, *Proc. R. Soc. A* **356**, 375 (1977).
- [211] F. Borgonovi, F. Izrailev, L. Santos, and V. Zelevinsky, “Quantum chaos and thermalization in isolated systems of interacting particles”, *Phys. Rep.* **626**, 1 (2016).
- [212] A. M. García-García and J. Wang, “Semi-Poisson statistics in quantum chaos”, *Phys. Rev. E* **73**, 036210 (2006).

- [213] N. Pancotti, G. Giudice, J. I. Cirac, J. P. Garrahan, and M. C. Bañuls, “Quantum East Model: Localization, Nonthermal Eigenstates, and Slow Dynamics”, *Phys. Rev. X* **10**, 021051 (2020).
- [214] L. Vidmar and M. Rigol, “Generalized Gibbs ensemble in integrable lattice models”, *J. Stat. Mech.: Theory Exp.* **2016** (6), 064007.
- [215] F. H. L. Essler and M. Fagotti, “Quench dynamics and relaxation in isolated integrable quantum spin chains”, *J. Stat. Mech.: Theory Exp.* (6), 064002.
- [216] J.-Y. Desaulles, A. Hudomal, C. J. Turner, and Z. Papić, “Proposal for Realizing Quantum Scars in the Tilted 1D Fermi-Hubbard Model”, *Phys. Rev. Lett.* **126**, 210601 (2021).
- [217] J. Haegeman, J. I. Cirac, T. J. Osborne, I. Pižorn, H. Verschelde, and F. Verstraete, “Time-Dependent Variational Principle for Quantum Lattices”, *Phys. Rev. Lett.* **107**, 070601 (2011).
- [218] A. A. Michailidis, C. J. Turner, Z. Papić, D. A. Abanin, and M. Serbyn, “Slow Quantum Thermalization and Many-Body Revivals from Mixed Phase Space”, *Phys. Rev. X* **10**, 011055 (2020).
- [219] V. Karle, M. Serbyn, and A. A. Michailidis, “Area-Law Entangled Eigenstates from Nullspaces of Local Hamiltonians”, *Phys. Rev. Lett.* **127**, 060602 (2021).
- [220] A. Hallam, *Embedding semiclassical periodic orbits into chaotic many-body Hamiltonians* (2023), at “Tensor networks for quantum many-body dynamics: from ergodicity breaking to dynamical phases of matter”.
- [221] P.-W. Lo and M. J. Lawler, “The fate of topological frustration in quantum spin ladders and generalizations”, *arXiv:2207.09549 [cond-mat.str-el]* (2022).
- [222] D. Perez-Garcia, F. Verstraete, M. M. Wolf, and J. I. Cirac, “Matrix Product State Representations”, *Quantum Inf. Comput.* **7**, 401 (2006).
- [223] A. Polkovnikov, K. Sengupta, A. Silva, and M. Vengalattore, “Colloquium: Nonequilibrium dynamics of closed interacting quantum systems”, *Rev. Mod. Phys.* **83**, 863 (2011).
- [224] M. Serbyn, Z. Papić, and D. A. Abanin, “Local conservation laws and the structure of the many-body localized states”, *Phys. Rev. Lett.* **111**, 127201 (2013).
- [225] L. Rademaker and M. Ortuño, “Explicit Local Integrals of Motion for the Many-Body Localized State”, *Phys. Rev. Lett.* **116**, 010404 (2016).
- [226] A. Y. Kitaev, “Unpaired Majorana fermions in quantum wires”, *Phys. Usp.* **44**, 131 (2001).
- [227] C. Dasgupta and S. K. Ma, “Low-temperature properties of the random Heisenberg antiferromagnetic chain”, *Phys. Rev. B* **22**, 1305 (1980).
- [228] D. S. Fisher, “Random transverse field Ising spin chains”, *Phys. Rev. Lett.* **69**, 534 (1992).
- [229] F. Iglói and C. Monthus, “Strong disorder RG approach of random systems”, *Phys. Rep.* **412**, 277 (2005).
- [230] A. C. Potter, R. Vasseur, and S. A. Parameswaran, “Universal Properties of Many-Body Delocalization Transitions”, *Phys. Rev. X* **5**, 031033 (2015).

- [231] J. F. Karcher, M. Sonner, and A. D. Mirlin, “Disorder and interaction in chiral chains: Majoranas versus complex fermions”, *Phys. Rev. B* **100**, 134207 (2019).
- [232] J. R. Schrieffer and P. A. Wolff, “Relation between the Anderson and Kondo Hamiltonians”, *Phys. Rev.* **149**, 491 (1966).
- [233] B. Ware, D. Abanin, and R. Vasseur, “Perturbative instability of nonergodic phases in non-Abelian quantum chains”, *Phys. Rev. B* **103**, 094203 (2021).
- [234] J. Dehaene and B. DeMoor, “Clifford group, stabilizer states, and linear and quadratic operations over GF(2)”, *Phys. Rev. A* **68**, 042318 (2003).
- [235] S. Aaronson and D. Gottesman, “Improved simulation of stabilizer circuits”, *Phys. Rev. A* **70**, 052328 (2004).
- [236] D. Gottesman, *Stabilizer codes and quantum error correction*, Ph.D. thesis, California Institute of Technology (1997).
- [237] G. Vidal, “Class of quantum many-body states that can be efficiently simulated”, *Phys. Rev. Lett.* **101**, 110501 (2008).
- [238] M. Serbyn, Z. Papić, and D. A. Abanin, “Criterion for Many-Body Localization-Delocalization Phase Transition”, *Phys. Rev. X* **5**, 041047 (2015).
- [239] M. Kiefer-Emmanouilidis, R. Unanyan, M. Fleischhauer, and J. Sirker, “Slow delocalization of particles in many-body localized phases”, *Phys. Rev. B* **103**, 024203 (2021).
- [240] T. Pető, F. Iglói, and I. A. Kovács, “Random Ising chain in transverse and longitudinal fields: Strong disorder RG study”, *Condensed Matter Physics* **26**, 13102 (2022).
- [241] M. Schechter, “LiHo_xY_{1-x}F₄ as a random-field Ising ferromagnet”, *Phys. Rev. B* **77**, 020401(R) (2008).
- [242] Y. Imry and S.-k. Ma, “Random-Field Instability of the Ordered State of Continuous Symmetry”, *Phys. Rev. Lett.* **35**, 1399 (1975).
- [243] J. Surace and L. Tagliacozzo, “Fermionic Gaussian states: an introduction to numerical approaches”, *SciPost Phys. Lect. Notes* , 54 (2022).
- [244] J. H. P. Colpa, “Diagonalisation of the quadratic fermion Hamiltonian with a linear part”, *J. Phys. A: Math. Gen.* **12**, 469 (1979).
- [245] A. Jafarizadeh and M. A. Rajabpour, “Entanglement entropy in quantum spin chains with broken parity number symmetry”, *SciPost Phys.* **12**, 195 (2022).
- [246] R. Alicki, M. Horodecki, P. Horodecki, and R. Horodecki, “On Thermal Stability of Topological Qubit in Kitaev’s 4D Model”, *Open Syst. Inf. Dyn.* **17**, 1 (2010).
- [247] F. Machado, D. V. Else, G. D. Kahanamoku-Meyer, C. Nayak, and N. Y. Yao, “Long-range prethermal phases of nonequilibrium matter”, *Phys. Rev. X* **10**, 011043 (2020).
- [248] D. V. Else, W. W. Ho, and P. T. Dumitrescu, “Long-lived interacting phases of matter protected by multiple time-translation symmetries in quasiperiodically driven systems”, *Phys. Rev. X* **10**, 021032 (2020).
- [249] D. E. Parker, R. Vasseur, and T. Scaffidi, “Topologically Protected Long Edge Coherence Times in Symmetry-Broken Phases”, *Phys. Rev. Lett.* **122**, 240605 (2019).

- [250] A. Chandran, A. Pal, C. R. Laumann, and A. Scardicchio, “Many-body localization beyond eigenstates in all dimensions”, *Phys. Rev. B* **94**, 144203 (2016).
- [251] M. Sbroscia, K. Viebahn, E. Carter, J.-C. Yu, A. Gaunt, and U. Schneider, “Observing Localization in a 2D Quasicrystalline Optical Lattice”, *Phys. Rev. Lett.* **125**, 200604 (2020).
- [252] C. Berke, E. Varvelis, S. Trebst, A. Altland, and D. P. DiVincenzo, “Transmon platform for quantum computing challenged by chaotic fluctuations”, *Nat. Commun.* **13**, 1 (2022).
- [253] A. Omran, H. Levine, A. Keesling, G. Semeghini, T. T. Wang, S. Ebadi, H. Bernien, A. S. Zibrov, H. Pichler, S. Choi, J. Cui, M. Rossignolo, P. Rembold, S. Montangero, T. Calarco, M. Endres, M. Greiner, V. Vuletić, and M. D. Lukin, “Generation and manipulation of Schrödinger cat states in Rydberg atom arrays”, *Science* **365**, 570 (2019).
- [254] E. Wybo, M. Knap, and F. Pollmann, “Entanglement dynamics of a many-body localized system coupled to a bath”, *Phys. Rev. B* **102**, 064304 (2020).
- [255] M. Ippoliti and W. W. Ho, “Solvable model of deep thermalization with distinct design times”, *Quantum* **6**, 886 (2022), 2208.10542v2 .
- [256] P. W. Claeys and A. Lamacraft, “Ergodic and Nonergodic Dual-Unitary Quantum Circuits with Arbitrary Local Hilbert Space Dimension”, *Phys. Rev. Lett.* **126**, 100603 (2021).