



Trinity College Dublin
Coláiste na Tríonóide, Baile Átha Cliath
The University of Dublin

Quantum Circuits

Transport, anomalous relaxation, and complexity

Alessandro Summer

School of Physics

Trinity College Dublin

A thesis submitted in partial fulfilment of the requirements towards the degree of
Doctor of Philosophy in Physics

Hilary Term 2026

Declaration

I declare that this thesis has not been submitted as an exercise for a degree at this or any other university and it is entirely my own work.

I agree to deposit this thesis in the University's open access institutional repository or allow the Library to do so on my behalf, subject to Irish Copyright Legislation and Trinity College Library conditions of use and acknowledgement.

I consent to the examiner retaining a copy of the thesis beyond the examining period, should they so wish (EU GDPR May 2018).

Signed: Alessandro Summer

Date: 24 February 2026

Quantum Circuits: Transport, anomalous relaxation, and complexity

Alessandro Summer

Hilary Term 2026

Abstract

Gate-based quantum computers can implement many-body system dynamics through repeated gate cycles. This periodic driving discretises time and relaxes energy conservation, while locality and symmetries still constrain *transport*, *relaxation*, and *complexity*. This thesis focuses on Floquet-circuit diagnostics linking spectral structure, transport, and operator complexity.

Part I introduces spectral and linear-response tools for Hamiltonian and Floquet systems, spectral projections of long-time limits, operator-resolved spectral densities (DOS/LDOS), and current spectral functions, and adapts them to stroboscopic dynamics with $U(1)$ conservation. We implement a moment-expansion (KPM) method for spectral densities as a hybrid quantum-classical workflow, and validate it by extracting the interacting-spin DOS on a trapped-ion quantum computer with up to $N=18$ spins. Using disordered discrete-time XXZ circuits with conserved charge, we benchmark ergodicity and transport. Because experiments access only finitely many cycles, we propose two finite-time transport estimators from Z -basis measurements: from the local-quench profile and its first circular moment we obtain drift and circular-variance measures that are more robust than conventional diagnostics. In this platform they reveal a long-lived near-SWAP prethermal *swappy* regime.

To probe discrete-time transport spectroscopically, we reconstruct spectral densities of stroboscopic correlators and derive a current operator valid for arbitrary circuit geometries. This yields quasienergy-resolved transport diagnostics, showing that a single integrability-breaking impurity can induce a prethermal plateau when studying quasi-conserved charges while leaving the current spectral function unchanged.

In parallel, we recast anomalous equilibration, including thermal and symmetry Mpemba effects, in resource-theoretic terms: monotone crossings are set by overlaps with slow, resourceful modes. The criterion extends to complexity-like resources and two-step Pontus-Mpemba protocols in unitary circuits, and the symmetry Mpemba effect can arise from a purely classical mechanism.

Part II studies operator growth via Krylov complexity in Floquet circuits. We relate continuous-time (Lanczos) and discrete-time (Arnoldi) recursion through a spectral-aliasing (“folding”) identity and show that defining an effective generator $\log \hat{U}$ is a physical branch choice. We control this choice with a locality-guided adiabatic unfolding protocol that tracks eigenprojectors along a Floquet path, producing quasi-local effective Hamiltonians when they exist. Benchmarks across dual-unitary endpoint families separate bounded free-fermion dynamics from qualitatively different interacting behaviour, clarifying when Krylov diagnostics require care for nonlocal Floquet generators.

Overall, the thesis provides experimentally native probes and reconstruction tools connecting discrete-time spectra, transport, and operator complexity to diagnose integrability breaking, prethermalisation, and anomalous relaxation in driven quantum matter.

List of Publications

The results presented in this thesis are based, primarily, on the following publications:

Published

1. Summer A., Chiaracane C., Mitchison M. T., Goold J.,
Calculating the Many-Body Density of States on a Digital Quantum Computer,
[Phys. Rev. Research 6, 013106 \(2024\)](#)
2. Summer A., Nico-Katz A., Dooley S., Goold J.,
Anomalous transport in $U(1)$ -symmetric quantum circuits,
[npj Quantum Information 12, 32 \(2026\)](#)
3. Summer A., Moroder M., Bettmann L. P., Turkeshi X., Marvian I., Goold J.,
Resource-theoretical unification of Mpemba effects: classical and quantum,
[Phys. Rev. X 16, 011065 \(2026\)](#)

Preprints

1. Aditya S., Summer A., Sierant P., Turkeshi X.,
Mpemba Effects in Quantum Complexity,
[arXiv:2509.22176 \(2025\)](#)

In preparation

1. Summer A., Dooley S., Goold J.,
Prethermalisation induced by an impurity in a $U(1)$ -symmetric quantum circuit,
Manuscript in preparation (2026)
2. Summer A., Yadalam H. K., Mitchison M.,
Krylov complexity of Floquet generators of quantum circuits,
Manuscript in preparation (2026)

The following are other publications of the author which, although relevant to a certain extent to this thesis and completed during the doctoral training, are not covered in detail in the following document:

1. Blekos K., Brand D., Ceschini A., Chou C.-H., Li R.-H., Pandya K., Summer A.,
A review on Quantum Approximate Optimization Algorithm and its variants,
[Phys. Rep. 1068, 1–66 \(2024\)](#)

Thanks

To John, for being a point of reference, sometimes a big brother, and someone who helped me think more clearly about the future. Over these four years it has been amazing to be drawn into your vortex of creativity and depth, and to see the originality you bring to everything, including supervision. You have been a constant inspiration.

To Mark, for remembering the best moments, the smallest details, and all the stupid things I was after saying along the way across four years of lunches, meetings, journal clubs, events, and pubs. Thanks for giving me the chance to have a go at expressing myself in scientific and non-scientific projects, and for always being so mindful about my development.

$1 + 1 > 2$. Thanks for teaching me how to better navigate the literature and my own head. I couldn't be luckier, and each of you taught me a great deal, in your own way, about physics and about life.

To Mattia, for always making me feel minded and looked after, and for sharing a fine appreciation of conspiracy theories and hours-long YouTube interviews. I cannot wait to see you become a father.

To Shane and Alex: those afternoons at DIAS, in front of a blackboard, helped me rediscover the excitement of research more than you probably realise. The freedom and creativity of those discussions meant more to me than any paper.

To Sindre, for your rare ability to listen properly. To Peter, for your sweetness.

To Hari, for your cheerful smile, which could change the mood even on the darkest days in the Dungeon. To Saulo and Khalak, for sharing that same underground solidarity.

To my QuSyster siblings, Leonard, Oisín, Laetitia, and Artur: it was a privilege to share this path with you. I wish you the very best for whatever comes next.

To my Dublin family: the warm, generous friends who made the place feel less like somewhere I was passing through and more like somewhere I belonged. To the Italians, Fra, So, Peanuts, Anita, the 830 polletti, Giò, Leo, Dani, and everyone else who made Dublin feel like home.

To my past, which has not stayed in the past, but has become an even stronger part of my present, the BB: Giamma, Gino, Ste, Fra, Scalo, Andre, Ila, and Simo. Distance has not stopped us from staying close, or from still managing to spend some of our best time together.

To my oldest friends, the “parish” ones, who have been there through all the different stages of our lives: Matteo, Gian, Leo, Chicco, Jay, Sara, Giulio.

To Emer, for giving me, in the end, the eyes of a local with which I rediscovered a city after four years, making me feel the transition that Ireland gave was finally complete. The QuSys legacy is in excellent hands.

To my family, as close and kind to me as I could ever hope for. Wherever I go, I carry the safety and confidence you have given me. Thank you, Grazia, my granny, for waiting almost until the end.

This thesis is submitted only in partial fulfilment of the PhD. The rest of the loot from these years is listed here.

Contents

Chapter 1: The role of Floquet circuits in quantum computation	1
1.1 Transport versus chaos: logically distinct questions	3
1.2 Thesis guide	4
 I Characterising transport	 7
Introduction to transport in many-body systems	8
Chapter 2: Infinite-time dynamics	15
2.1 Spectral projections.	17
2.1.1 Eigenvalue distribution	19
2.1.2 Eigenstate thermalisation and long-time transport	20
2.1.3 Structure of eigenstates	21
2.2 Spectral densities	23
2.3 Currents and Drude weight	25
2.4 Spectral densities of a XYZ chain.	27
2.4.1 Continuous time	28
2.4.2 Discrete time	29
2.5 Integrable discrete-time XXZ chain.	32
2.5.1 Dual-unitary circuits.	34
2.6 Discrete-time disordered XXZ chain	35
2.6.1 Overview of regimes of transport dynamics.	37
2.6.2 Spectral properties of a discrete-time model	38
2.7 Current operators and the single-impurity circuit	40
2.7.1 Continuous time	41
2.7.2 Discrete time	43
2.7.3 Numerical protocol and results	45
2.8 Conclusion	48
2.9 Methods	49
2.9.1 POLFED	50
2.9.2 Efficient computation of the entropy in a magnetisation subsector	50
2.9.3 Stochastic trace evaluation and recursion	51
2.9.4 The Kernel Polynomial Method	52
2.9.5 Kernel polynomial method as a quantum algorithm	55
2.9.6 Implementation of the KPM on Quantinuum H1-1	59
2.9.7 Spectral reconstructions for discrete-time circuits	62
Chapter 3: Finite-time dynamics	65
3.1 Floquet prethermalisation as a finite-time window	66
3.1.1 High-frequency prethermalisation and a dressed generator $\hat{H}_*$. .	67
3.1.2 Prethermalisation near an integrable Floquet point	67
3.2 Transport diagnostics from magnetisation profile.	68
3.2.1 Local magnetisation profile in a fixed- M sector	69
3.2.2 Semiprobability profile in the $M=0$ sector	69
3.2.3 From spreading profiles to hydrodynamic transport	70
3.2.4 Circular ensemble and its relation to the dynamic structure factor	70

3.2.5	Wrapped-normal reference and effective exponents	72
3.3	Disordered Floquet XXZ circuit	73
3.4	Prethermalisation near the SWAP point: the swappy regime	74
3.5	Transport observables from drift, spreading, and peak decay	76
3.6	Extracting effective exponents and drift velocities	77
3.7	Single-impurity XXZ circuit and prethermalisation	78
3.7.1	Conserved charges of the clean XXZ circuit	79
3.7.2	Impurity Floquet step and autocorrelator probe	80
3.7.3	Leading resonance interpretation and the role of the Heisenberg time	82
3.8	Conclusion	82
3.9	Methods	83
3.9.1	Efficient time evolution of a state by a random circuit	83
3.9.2	Numerical methods: vector propagation and typicality in the $M=0$ sector	84
Chapter 4:	Anomalous equilibration: Mpemba effect	87
4.1	Details on the Davies map	89
4.2	Resource Theories	90
4.3	The thermal Mpemba effect	91
4.3.1	Athermality as a Resource	92
4.3.2	Example 1: Thermalisation of a Spin Chain	93
4.3.3	Example 2: The Davies Map	96
4.3.4	Example 3: Thermal Mpemba Effect with ETH	98
4.4	The symmetry Mpemba effect	99
4.4.1	Symmetries in Quantum Systems	100
4.4.2	Asymmetry as a Resource	104
4.4.3	Modes of Asymmetry	104
4.4.4	Example 1: A Classical Symmetry Mpemba Effect with $\mathbb{Z}_4$. . .	107
4.4.5	Example 2: A Quantum Symmetry Mpemba Effect with $\mathbb{Z}_4$. . .	109
4.4.6	Example 3: Time-Translation-Symmetry Mpemba Effect in Davies Maps	111
4.4.7	Example 4: Symmetry Mpemba Effects in Quantum circuits . .	112
4.5	A unified description of the thermal and the symmetry Mpemba effects	118
4.6	Complexity Mpemba effects	120
4.6.1	Resource of imaginarity	122
4.6.2	Resource of non-Gaussianity	123
4.6.3	Resource of magic	124
4.6.4	Mpemba effects of complexity	125
4.6.5	Pontus-Mpemba Effects	126
4.7	Conclusion	128
II	Characterising complexity	129
	Complexity, chaos, and transport	130
Chapter 5:	Krylov Complexity	135
5.1	Introduction	135
5.2	Krylov bases: continuous vs. discrete time	137
5.2.1	Continuous time: Lanczos	139
5.2.2	Discrete time: Arnoldi	140

5.3	Connect Lanczos and Arnoldi	143
5.3.1	Constant Lanczos coefficients	146
5.3.2	MEU to Lanczos	148
5.4	Floquet Hamiltonians	150
5.5	Maximum Eigenprojector Overlap	153
5.5.1	Floquet path and tracking	153
5.5.2	Eigenprojectors near degeneracies	154
5.5.3	Assignment via projector overlap	155
5.5.4	Quasienergy unfolding	155
5.5.5	Locality tie-breaker	156
5.6	XX spin chain under brickwork drive	156
5.6.1	Analytic phases and unfolding	158
5.6.2	Comparison of unfolding methods	160
5.6.3	XX-DU point Krylov complexity	162
5.7	Other DU models	163
5.8	Conclusion	166
Chapter 6: Summary and outlook		169
6.1	Outlook	172
6.2	Closing perspective	173
Appendix		175
A	Connection to quantum singular value transformation	176
A.1	QSVT in brief	176
A.2	Chebyshev polynomials from the signal unitary	176
A.3	Spectral-bulk arccosine approximation as a NISQ surrogate	177
A.4	What is preserved and what is traded away	178
B	U(1) two-qubit gate parametrisations	179
B.1	From four local Z -phases to (h_n, h'_n, ϕ_n)	179
B.2	Explicit 4×4 matrix of $\hat{u}_{n,n+1}$	180
B.3	Mapping to the Žnidarič-Duh-Zadnik (ZDZ) parametrisations	180
B.4	Simulation-friendly decompositions	182
C	Transport protocol technicalities	183
C.1	Dynamical correlation function in the full Hilbert space	183
C.2	Restriction to the fixed- M subspace	183
C.3	Typicality within $\mathcal{H}^{(M)}$ and the $M = 0$ semiprobability	184
D	Spectral properties and light-cone correlations from a transfer matrix	184
E	From global G -invariance to local G -covariance	185
F	Asymptotic decay of the quantum relative entropy	186
F.1	Proof of the Asymptotic Decay of the KL-Divergence	187
F.2	Proof of the Asymptotic Decay of the Quantum Relative Entropy	188
G	Self-dual kicked Ising gate	189
G.1	Properties of the SDKI two-site gate	189
G.2	Connection to the four-term logarithm	191
G.3	Connection to the principal six-term logarithm	192
List of Acronyms		196
References		197

Chapter 1

The role of Floquet circuits in quantum computation

Physics is once again being reshaped by a resurgence of quantum ideas, not because the postulates of quantum mechanics have changed, but because our ability to deploy them as laboratory tools has. Roughly a century after quantisation, superposition, and the uncertainty principle reshaped the foundations of physics, the field is now undergoing a technological transformation often labelled the “second quantum revolution” [8]. The term is not particularly precise, but it captures a useful distinction. The first wave of quantum-enabled technologies translated microscopic quantum effects into reliable macroscopic functions; semiconductors, transistors, lasers, and magnetic resonance techniques are canonical examples. These devices exploit quantum mechanics in regimes where dissipation and material complexity yield robust classical behaviour.

Second
quantum
revolution

The modern wave is different in its target: not the passive use of quantum effects, but the ability to exercise fine-grained control over them. Preparing, manipulating, and measuring *individual* quantum systems with long-lived coherence has become routine, since the first experiments of the 1980s. Trapped ions, neutral atoms, and superconducting circuits now allow control of many-qubit dynamics [9–12]. The focus here is primarily not on scale, but on the ability of fine control: rather than relying on ensemble-averaged effects, contemporary platforms deliberately engineer superposition, entanglement, and measurement backaction as programmable resources [8, 13].

Fine-
grained
control

The applications emerging from the second revolution range in maturity. Quantum sensing is now a well-developed paradigm for surpassing classical precision bounds in specialised metrological tasks [14], and quantum cryptography is technologically established, supported by an extensive theoretical literature and repeated real-world deployments [15]. As for quantum computing, although it remains comparatively early-stage, it represents the most conceptually radical advancement: information encoded in qubits can be processed through controlled interference among exponentially many computational amplitudes, enabling dramatic speedups for specific problems such as integer factoring and, more broadly, the simulation of quantum matter [16–19]. A computational model is “digital” only when the acceptable input-output state space is countable and the allowed transformations can be implemented via a discrete set of

Applications

transition steps specified by a discrete set of parameters; otherwise it is “analogue.” In this sense, present-day noisy intermediate-scale quantum (NISQ) processors, despite being programmable, often remain essentially analogue because full fault-tolerant digitisation of quantum information is not yet available [20]. Recent large-scale, digital quantum computer demonstrations reinforce that also this side of the second revolution is increasingly becoming more concrete [13, 21].

Why
Floquet?

Researchers choose what to focus on not only for its long-horizon utility, but also for its tractability and experimental feasibility. In practice, this favours models that are analytically transparent, numerically benchmarkable, and experimentally realisable. The harmonic oscillator plays this role in equilibrium physics because it yields exact solutions and anchors perturbative expansions around stable baselines across disciplines [22]. For digital quantum devices, the analogous baseline is not a single Hamiltonian but a *protocol*: a time-dependent drive envelope or a discrete update rule. Floquet circuits, understood broadly as periodically driven systems, whether from time-dependent Hamiltonians or from stroboscopic circuits, provide a systematic language for such protocols and have become standard workhorses for nonequilibrium physics and control [23, 24].

Classical
analogy

This perspective becomes particularly natural when one looks at how present-day hardware is used. Classical computation is physically realised through semiconductor devices whose collective switching can be abstracted as logic gates; at the lowest algorithmic level, iterating local update rules across space and time yields cellular automata, where complex emergent behaviour arises from simple local maps [25, 26]. Traffic-like dynamics in elementary cellular automata such as rule 184 illustrates how rich phenomenology can emerge from strictly local deterministic updates [27, 28].

Circuit
viewpoint

Digital quantum processors are controlled in a closely analogous way. Their elementary primitives are calibrated pulses, waveforms, or engineered couplings that realise local operations (quantum channels). A full cycle of such operations defines a discrete-time propagator, and repeating the cycle yields the many-body evolution $\hat{U}_{\text{cyc}}^t$. In this sense, local quantum circuits are the natural stroboscopic counterparts of classical cellular automata: both are built from parallel local update rules, while the quantum version preserves coherence and generates entanglement [13, 29]. Quantum cellular automata, in particular, formalise this connection by imposing translation invariance and strict causality on the global update, and can be viewed as a structured subclass of Floquet circuit dynamics¹ [30–32].

¹Quantum cellular automata are often defined as translation-invariant, causal, discrete-time dynamics on a lattice; within that framework, the familiar brickwork (Margolus) partitioning appears as a canonical circuit representation [31, 32].

Viewed experimentally, running a quantum processor at the circuit (Floquet) level is first and foremost an experimental protocol: the device realises a concrete physical process whose input-output statistics can be measured, characterised, and compared to an idealised model. When the implemented experimental protocol $\hat{U}_{\text{phys}}$ is close to the intended gate model $\hat{U}$ in an appropriate measurable distance (e.g. a diamond-norm or process-infidelity metric), the matrix description is predictive; when it is not, the observed evolution defines a distinct empirical object that must itself be modelled via noise channels, calibration drift, and effective generators [29, 33]. This point is not conceptual but methodological: discrete-time quantum evolution is experimentally primary in gate-based platforms, and effective Floquet descriptions are part of the basic toolkit for extracting physics from hardware.

Experimental
viewpoint

This convergence between experimental setup and theoretical language has turned discrete-time quantum dynamics into a central laboratory for nonequilibrium many-body physics. Periodic driving and circuit dynamics can stabilise long-lived prethermal regimes, generate effective Hamiltonians unavailable in static settings, and produce intrinsically nonequilibrium phases such as discrete time crystals and anomalous Floquet topological matter [23, 24, 34–37]. At the same time, local circuits have become a standard setting for quantitative theories of information spreading, entanglement growth, and hydrodynamic relaxation, and they host genuinely new kinds of dynamical criticality when unitary evolution competes with measurements [38–40]. Even within strictly unitary dynamics, special circuit classes such as dual unitary (DU) models provide rare examples combining exact solvability with strongly mixing, chaotic-looking behaviour, enabling unusually sharp statements about correlations and transport [41].

Floquet
laboratory

This chapter provides the conceptual entry point. [Sec. 1.1](#) first fixes the minimal terminology (transport, chaos, ergodicity, scrambling) used throughout the thesis. [Sec. 1.2](#) then gives a chapter-by-chapter guide to [Part I](#) and [Part II](#).

1.1 Transport versus chaos: logically distinct questions

The words *transport*, *chaos*, and *scrambling* are sometimes used interchangeably in the many-body literature, but in this thesis they answer different questions. [Part I](#) develops transport probes, while [Part II](#) develops operator-complexity diagnostics; here we only fix the minimal terminology needed to differentiate among these threads.

Transport
(infrared)

Transport is about the spatio-temporal relaxation of *conserved* densities at long wavelengths and late times. In practice this means hydrodynamic scaling, exponents, and (when appropriate) coefficients such as diffusion constants, extracted from currents and density profiles [42, 43].

Chaos
versus
ergodicity

Chaos is a statement about dynamical genericity once symmetries are resolved; in quantum many-body systems we probe it operationally through spectral universality (random matrix theory (RMT) statistics [44]) and related eigenstate and response signatures rather than through literal trajectory divergence [45, 46]. *Ergodicity* is a statement about long-time exploration/typicality: for generic initial conditions, local observables relax to the maximum-entropy ensemble compatible with the conserved quantities (in Floquet settings, typically the infinite-temperature ensemble within each symmetry sector); eigenstate thermalisation hypothesis (ETH) is a common sufficient mechanism for this thermalisation logic in Hamiltonian systems [46]. The two are correlated but not identical: one can have random-matrix-like spectra yet atypical eigenstates or slow nonthermal subspaces (e.g. scars), and conversely the onset of thermal behaviour need not uniquely fix spectral correlations [47].

Scrambling
(operator
spreading)

Scrambling refers to how rapidly initially local quantum information delocalises under time evolution. It is naturally probed through operator growth diagnostics such as out-of-time-order correlators (OTOCs) and associated light-cone/butterfly scales [39, 48, 49]. Scrambling rates need not track transport coefficients, and conservation laws can slow operator growth even when spectral diagnostics look chaotic.

In Floquet settings this separation is especially clean: the Floquet unitary can look spectrally chaotic while charge transport ranges from ballistic to diffusive to anomalous depending on constraints and integrable structure.

1.2 Thesis guide

The rest of the thesis is structured around two complementary goals: (i) diagnosing transport and equilibration in U(1)-symmetric Floquet chains using probes that are meaningful in discrete time, and (ii) quantifying how operator structure grows once one goes beyond purely hydrodynamic questions.

Part I:
transport

In Ch. 2 we examine an infinite-time (spectral) viewpoint: time averages as spectral projections, operator-resolved spectral measures (density of states (DOS), local density of states (LDOS)), and spectral functions that connect directly to linear response and to static chaos diagnostics. This chapter reports the results of Refs. [1, 2].




Figure 1.1: Overview of the thesis structure. We differentiate between chaos, transport, relaxation, and complexity diagnostics, and between infinite-time (Ch. 2) and finite-time (Ch. 3) probes. Ch. 4 focuses on state-dependent relaxation rates, while Ch. 5 develops operator-growth diagnostics. The thesis is organised into Part I (Chs. 2 to 4), centred on transport and relaxation, and Part II (Ch. 5). Although Part II highlights how these threads can be viewed as all complexity probes, Ch. 5 is the only chapter whose primary focus is complexity.

In Ch. 3 we turn to finite-time dynamics, as Floquet experiments and numerics are fundamentally stroboscopic and often dominated by long-lived intermediate-time regimes (e.g. crossover and prethermal regimes). We quantify transport via local-quench profile diagnostics and complementary current-correlation viewpoints, and identify Floquet-specific prethermal behaviour (the “swappy” regime of Ref. [2]) that is invisible to infinite-time diagnostics alone.

In Ch. 4, reviewing the results of Refs. [3, 4], we shift the emphasis from transport coefficients to how relaxation proceeds across different initial conditions. Using a resource-theoretic framing, it explains anomalous equilibration (Mpemba effects) in terms of overlaps with resourceful slow modes, providing a unifying language for state-dependent relaxation rates that will be useful again when discussing operator growth.

In Part II we shift focus to operator growth as a complexity diagnostic, logically distinct from both spectral chaos and transport. Ch. 5 develops Krylov complexity as a seed-resolved embedding of operator dynamics in a one-dimensional Krylov chain, placing continuous-time (Lanczos) and discrete-time Floquet (Arnoldi) constructions on the same footing. It derives the spectral-aliasing relation that connects the associated measures and recursion data, and addresses the branch ambiguity of $\log \hat{U}$ via a locality-guided adiabatic unfolding protocol (maximum eigenprojector overlap (MEO)), benchmarked in driven XX and dual-unitary circuit families.

Part II:
complexity

As a reading aid, as described in Fig. 1.2, the header bar is anchored at the spine and lets the reader track both the current chapter or section and the current page position,

globally across the thesis body on left pages and locally within the active chapter on right pages.




Figure 1.2: The header bar is anchored at the spine. On left-hand (verso) pages it shows the current chapter within the thesis body (Chs. 1 to 6); because the inner margin is on the right, the page order is mirrored and the indicator advances from right to left. On right-hand (recto) pages it shows the current section within the active chapter; there the inner margin is on the left, so the indicator advances from left to right. In both cases, the dot marks the current page position, measured over the thesis body on left pages and over the current chapter on right pages.

Part I

Characterising transport

Introduction to transport in many-body systems

Transport in this thesis refers to the spreading of conserved quantities in interacting quantum many-body systems, with a primary focus on U(1) charge (magnetisation) transport in spin chains and U(1)-symmetric Floquet brickwork circuits. The emphasis on conserved densities is deliberate: conservation laws guarantee slow, long-wavelength modes and hence meaningful transport coefficients in the thermodynamic limit [42, 43].

We will quantify transport using two complementary diagnostics, aligned with the two time-scales emphasised in this thesis.

Infinite-
time
transport

(i) *Infinite-time (spectral) transport*, Ch. 2. Here the focus is on late-time limits and frequency-domain objects determined by the spectrum (eigenvalues or eigenphases) and eigenvectors of the generator of dynamics (either $\hat{H}$ or the Floquet unitary $\hat{U}$), together with operator-resolved spectral measures such as the density of states (DOS) and local density of states (LDOS), and dynamical spectral functions. Transport coefficients in linear response are encoded in the current spectral function $S_J(\omega)$ in Lehmann (Kubo) form: its low-frequency structure separates a possible Drude contribution from a regular continuum [42, 43, 50].

Finite-
time
transport

(ii) *Finite-time (dynamical) transport*, Ch. 3. Here the focus is on experimentally and numerically accessible transients, quantified directly in the time domain. Concretely, we analyse the spreading of an initially localised U(1) charge excitation (spin flip) through the time-evolving magnetisation profile, extracting widths/second moments and scaling exponents. These profile diagnostics are closely connected to two-point functions of the conserved density (e.g. spin-spin correlators) and can be cross-checked against the current autocorrelator $C_J(t)$, which probes transport via the same current operator that enters linear response [42, 43, 50].

Because the main platform of this thesis is discrete-time (Floquet) dynamics, we use continuous-time Hamiltonian chains as benchmarks but ultimately formulate all diagnostics in a stroboscopic language: in the Floquet setting the generator is the one-period unitary $\hat{U}$, and both spectral objects (quasienergies/eigenphases and operator matrix elements) and time-domain probes are naturally sampled at integer numbers of periods.

Transport
taxonomy

A useful starting point is the single-particle diffusion paradigm. For Brownian motion, coarse-graining leads to the diffusion equation and the Einstein relation for the

mean-squared displacement (MSD) [51],

$$\langle x^2(t) \rangle = 2Dt \quad (\text{in } 1D), \quad (1.1)$$

which provides a mathematical definition of the diffusion constant D .

In many-body settings one typically extracts a *transport exponent* from the scaling of the width of an evolving density profile. For an initially localised excess density $\delta\rho(x, 0)$, it is natural to quantify the spreading by the *second cumulant* (variance), i.e. the central second moment. Define the packet centre

$$\bar{x}(t) = \frac{\int dx x \delta\rho(x, t)}{\int dx \delta\rho(x, t)}, \quad (1.2)$$

and the width

$$m_2(t) = \frac{\int dx (x - \bar{x}(t))^2 \delta\rho(x, t)}{\int dx \delta\rho(x, t)} \quad (\text{continuum}). \quad (1.3)$$

(or the corresponding discrete sum on a chain), and set $m_2(t) \sim t^\alpha$. The exponent α classifies transport: $\alpha = 2$ (ballistic), $\alpha = 1$ (diffusive), and $0 < \alpha < 2$ with $\alpha \neq 1$ (anomalous), with a broad range of anomalous behaviours possible in low-dimensional quantum systems [43]. Only in the *diffusive* case does it make sense to define a diffusion *coefficient* from late times via

$$D = \lim_{t \rightarrow \infty} \frac{m_2(t)}{2t} \quad (\text{in } 1D), \quad (1.4)$$

provided the asymptotic growth is linear.

The single-particle diffusion paradigm is conceptually useful, but in interacting many-body systems transport is typically *not* understood by following the trajectory of a tagged microscopic constituent. The reason is not merely that the dynamics is complicated, but that the experimentally accessible long-time, long-wavelength behaviour is controlled by *collective* slow variables: conserved densities and their associated currents [42, 43]. This motivates the continuum viewpoint pioneered in Fourier's phenomenology, where one relates currents to gradients and closes the resulting continuity equations by constitutive relations (diffusion being the simplest example) [52, 53].

Between microscopic dynamics and this hydrodynamic description there is an important intermediate possibility: in weakly correlated regimes the low-energy excitations can be reorganised into long-lived *quasiparticles*, so that transport may be described kinetically in terms of their propagation and scattering (e.g. phonons for heat in insulators, magnons/spinons for spin transport in certain magnets). When such quasiparticles are sharply defined, transport coefficients can often be estimated from their velocities

and mean free paths. However, in strongly interacting regimes quasiparticles may be short-lived or absent, and then transport is most naturally formulated directly at the level of conserved densities and currents, i.e. in hydrodynamic terms [42, 43].

Hydrody-
namics

Hydrodynamics is an emergent long-wavelength description of interacting many-body dynamics in one dimension. Its basic assumption is that, beyond microscopic relaxation times, local observables are well described by a *local equilibrium* state whose parameters (temperature and the chemical potential conjugate to a conserved density) vary slowly in space and time [42, 43]. The effective theory is then written in terms of coarse-grained fields (equivalently, local-equilibrium expectation values). In this thesis the conserved U(1) charge is the local magnetisation. We denote by $m(x, t)$ the coarse-grained magnetisation density and by $j(x, t)$ the associated current in the continuum description. For Hamiltonian dynamics, conservation implies the 1D continuity equation

$$\partial_t m(x, t) + \partial_x j(x, t) = 0. \quad (1.5)$$

For Floquet dynamics (one period set to unity), the analogous statement is a lattice continuity equation at stroboscopic times $t \in \mathbb{Z}$,

$$m_n(t+1) - m_n(t) + j_{n+1/2}(t) - j_{n-1/2}(t) = 0, \quad (1.6)$$

which reduces to Eq. (1.5) upon coarse-graining in space and on long stroboscopic times.

To close the theory one postulates a *constitutive relation*: an expansion of j in gradients of the local thermodynamic fields. For a U(1) charge the natural field is the local chemical potential $\mu(x, t)$ conjugate to $\hat{S}^z = \sum_n \hat{s}_n^z$, and to leading order one writes

$$j(x, t) = -\sigma \partial_x \mu(x, t) + \dots, \quad (1.7)$$

where σ is the dc conductivity (per unit length) and the ellipsis denotes higher-gradient and nonlinear corrections. Linearising local equilibrium gives $\delta m = \chi \delta \mu$ with χ the static susceptibility, hence Eq. (1.7) becomes Fick's law,

$$j(x, t) = -D \partial_x m(x, t) + \dots, \quad D = \frac{\sigma}{\chi}, \quad (1.8)$$

i.e. the Nernst-Einstein relation in one dimension. In a diffusive regime this implies broadening of a localised packet with variance $m_2(t) \sim 2Dt$ (with t either continuous time or Floquet periods, depending on the dynamics), providing a mathematical definition of D from real-space spreading.

Transport coefficients can be defined without preparing gradients or non-equilibrium initial states. Linear-response theory says that the response of a conserved density to a weak bias is fixed by *equilibrium* current fluctuations of the unperturbed dynamics [42, 50, 54]. This is the conceptual bridge between the two diagnostics used throughout [Part I](#): the same current operator that controls the spreading of density packets at finite times also controls the conductivity extracted from equilibrium correlations. Concretely, we focus on the total magnetisation current $\hat{J}^z = \sum_n \hat{j}_{n+1/2}^z$ and its equilibrium autocorrelator $C_J(t) = N^{-1} \langle \hat{J}^z(t) \hat{J}^z \rangle$. For Floquet circuits, t is an integer period index and all correlators are naturally sampled stroboscopically.

The Green-Kubo statement is then simply that the dc conductivity is governed by the time-integrated current fluctuations (in Floquet: the corresponding infinite-time *sum* over periods), and hydrodynamics relates this to diffusion via

$$D = \frac{\sigma_{\text{dc}}}{\chi}. \quad (1.9)$$

In practice, [Part I](#) uses this relation in reverse: we compute or reconstruct equilibrium correlation data (often in the frequency/spectral language of [Ch. 2](#)) and compare it against real-space spreading diagnostics ([Ch. 3](#)) whenever a diffusive interpretation is appropriate.

In a genuinely diffusive regime this gives a nontrivial consistency check: the D extracted from profile broadening and the D inferred from equilibrium current fluctuations must agree (up to finite-size and finite-time effects).

Equilibrium linear response also makes clear *what* can prevent diffusion. If the low-frequency conductivity contains a ballistic channel, it appears as a Drude peak,

$$\text{Re } \sigma(\omega) = 2\pi D_{\text{Drude}} \delta(\omega) + \sigma_{\text{reg}}(\omega). \quad (1.10)$$

Equivalently, ballistic transport corresponds to a non-decaying long-time component of the total-current autocorrelator (in Floquet: a nonzero long-time *stroboscopic* average).

A rigorous sufficient criterion for such a non-decaying component is the Mazur inequality (and Mazur-Suzuki refinements) [55, 56]. If $\{\hat{Q}_n\}$ is any set of (quasi-)local conserved operators with nonzero overlap with $\hat{J}^z$, then the long-time averaged autocorrelator obeys

$$C_\infty \geq \sum_n \frac{\langle \hat{J}^z \hat{Q}_n \rangle^2}{\langle \hat{Q}_n^2 \rangle}, \quad (1.11)$$

where C_∞ denotes the appropriate infinite-time average (also known as continuous-time Cesàro mean for Hamiltonians, or the period-averaged stroboscopic mean for Floquet dynamics). Thus, conserved quantities (including those enforced by

integrability or by circuit structure) can *force* a strictly ballistic contribution. What Mazur does *not* determine is the behaviour of $\sigma_{\text{reg}}(\omega)$, where diffusion and anomalous transport live. At finite size, boundary conditions affect how the ballistic weight appears: with periodic boundary condition (PBC) it is naturally tied to the $\omega = 0$ component, while with open boundary condition (OBC) it is typically redistributed into peaks at $\omega \sim 1/N$ so the robust diagnostic is the *integrated* low-frequency weight (see Sec. 2.3).

A further subtlety is that in isolated quantum many-body systems the microscopic dynamics is strictly unitary and the global pure state remains pure, yet local observables can relax and exhibit transport behaviour consistent with hydrodynamics [46, 57, 58]. It is therefore useful to distinguish *equilibration* from *thermalisation* in the ETH sense [46, 59].

Equilibration and
Ch. 4

Equilibration means that expectation values of *local* observables (equivalently, reduced density matrices of small subsystems) become approximately time-independent for most late times, approaching the values predicted by the *diagonal ensemble* associated with the initial state [57]. This is a dephasing phenomenon: writing $|\psi(0)\rangle = \sum_i c_i |E_i\rangle$, unitary evolution retains the phases $e^{-iE_i t}$, but for generic spectra the off-diagonal contributions to local observables dephase at late times, leaving a steady value controlled by $|c_i|^2$ rather than by the microscopic phases [46, 57]. One may view this as the initial *coherence* between energy shells becoming inaccessible to local probes as it is redistributed into increasingly nonlocal correlations, even though the full state remains pure [46, 57]. This is connected with the concept of quantum typicality that we will employ later [60]. The subsystem notion of equilibration above is also the one relevant for the Mpemba effects discussed in Ch. 4, where one compares rates of approach towards a specified reference manifold for the reduced state, rather than invoking equilibration of a globally isolated pure state.

Thermalisation

Thermalisation is the stronger statement that these equilibrated late-time local values coincide with those of an equilibrium ensemble fixed only by the *relevant conserved quantities* (for Hamiltonian dynamics: typically energy and global charges; for Floquet dynamics: typically only global charges) [46, 59]. In generic non-integrable systems this is explained by the ETH: diagonal matrix elements of few-body observables are smooth functions of energy and off-diagonal elements are exponentially small in system size, so that the diagonal ensemble becomes locally indistinguishable from the microcanonical (or Gibbs) ensemble at the corresponding energy density [46, 59]. Recent progress has also established rigorous conditions under which thermalisation must occur in translation-invariant qubit systems at high effective temperature (assuming the absence of perfect spectral resonances): typical low-complexity pure

states drawn from high-entropy ensembles become locally indistinguishable from a Gibbs state under unitary time evolution [61].

This distinction is based on the classical idea of *ergodicity*: for a measure-preserving flow Φ_t on phase space ergodicity denotes that the infinite-time average of an observable is equal to its phase-space (ensemble) average for nearly all initial conditions. This is exact for an integrable phase function f according to Birkhoff’s ergodic theorem.

$$\lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T f(\Phi_t(x)) dt = \langle f \rangle \quad \text{for almost every } x, \quad (1.12)$$

where the invariant-measure average is indicated by $\langle f \rangle$. Ergodicity breaking refers to strong violations of this behaviour while ergodic is used in an analogous sense in the many-body quantum literature to denote systems where typical highly excited eigenstates satisfy ETH and local observables thermalise (e.g. due to many-body localisation (MBL) [46, 58] or integrability). Recent work has promoted a complementary notion of *Hilbert-space ergodicity* for driven and controlled dynamics analysing both obstacles and design principles in Floquet settings [62] going beyond this traditional usage (ETH $\Rightarrow$ thermalisation of local observables). In a similar vein universal features of measured bitstring distributions have been proposed as experimentally accessible signatures of Hilbert-space ergodicity which can be used to diagnose and learn noise in quantum devices [63].

A large and well-studied class of non-ergodic models is the class of integrable models. An *integrable* many-body system is characterised (in the present context) by the existence of an extensive family of independent (quasi-)local conserved operators $\{\hat{Q}_n\}$ that commute with the dynamics generator and with each other,

$$[\hat{H}, \hat{Q}_n] = 0, \quad [\hat{Q}_n, \hat{Q}_m] = 0, \quad (1.13)$$

(or the Floquet analogue $[\hat{U}, \hat{Q}_n] = 0$), beyond the “trivial” conservation laws (energy, total U(1) charge, etc.) [43, 46]. These additional constraints obstruct generic ETH behaviour and lead to relaxation towards a generalised Gibbs ensemble (GGE) that retains memory of the full set $\{\hat{Q}_n\}$ [46]. At the level of transport, integrability necessitates an extension of standard hydrodynamics: generalised hydrodynamics (GHD) provides an effective description of large-scale non-equilibrium evolution in such settings and has been confronted with cold-atom experiments [64–68]. A distinct mechanism for failure of thermalisation is MBL, where strong disorder stabilises an extensive set of quasi-local integrals of motion and prevents thermal transport and ergodicity even without conventional integrability [58]. In particular, an extensive set of conserved charges makes it natural for the Mazur-Suzuki bound Eq. (1.11) to

Ergodicity

Integrability

yield a finite lower bound on the Drude weight whenever the spin current has nonzero overlap with some of these charges [43, 55, 56].

In generic non-integrable systems, by contrast, only a small set of global conservation laws typically survives, and long-time dynamics is often consistent with thermalisation and hydrodynamic transport described by a finite diffusion constant [43, 46]. Important exceptions arise in disordered systems (e.g. MBL), and in periodically driven systems where heating can qualitatively alter late-time behaviour; throughout this thesis we therefore keep the conserved $U(1)$ charge as the common reference point for comparing Hamiltonian and Floquet settings [46, 58].

Chapter 2

Infinite-time dynamics

A recurring lesson of statistical mechanics is that, in generic thermalising many-body systems, infinite-time long-wavelength dynamics is hydrodynamic and controlled by conserved densities. Local stationary properties are fixed by the equilibrium ensemble constrained by those conservation laws, with transport encoded in a small number of coefficients [42, 43, 69]. Already for classical ergodic systems one replaces microscopic trajectories by time-averaged observables; in quantum mechanics this logic sharpens because infinite-time limits are controlled by the spectral structure of the generator of dynamics. Building on the taxonomy of Part I, we now formalise this spectral (infinite-time) viewpoint.

In continuous time the generator is the Hamiltonian $\hat{H}$; in discrete time, the native language of gate-based devices and Floquet engineering, it is the one-period unitary $\hat{U}$. In both settings, infinite-time equilibration and transport can be posed as spectral questions: which frequencies (energies or eigenphases) occur, how spectral weight is distributed, and how strongly physically relevant operators couple to the corresponding spectral components.

The goal of this chapter is to develop that viewpoint in a form adapted to transport and ergodicity in Hamiltonian and Floquet circuits. We focus on three linked layers. First, *transport*: linear response relates transport coefficients to current-current correlators (Green-Kubo and the low-frequency structure of $S_J(\omega)$), while real-space spreading probes the same physics through the evolution of a localised perturbation. In the numerical sections we primarily use spin-profile broadening to distinguish ballistic, diffusive, and anomalous regimes, while correlator-based quantities enter naturally through the spectral objects and Drude-weight diagnostics [42, 43, 50, 70].

Second, *ergodicity*: eigenvalues (or eigenphases) and eigenvectors provide static chaos diagnostics, including level correlations (e.g. gap ratios and the spectral form factor (SFF)) and eigenstate structure (entanglement, participation, and ETH scaling of matrix elements), separating integrable and localised behaviour from random-matrix-like regimes [46, 59, 71–73]. Third, *operator coupling*: we use the DOS $g(E)$ and the LDOS (operator-resolved spectral density), together with related spectral functions, to organise how local operators decompose over eigenstates. These objects directly control overlaps appearing in Mazur-type bounds and relaxation rates, and therefore mediate between spectral structure and transport [55, 56, 74].

For Floquet systems the translation is systematic: energies are replaced by eigenphases, infinite-time averages become phase-projected (diagonal) ensembles of $\hat{U}$, and linear-response formulas admit stroboscopic analogues in which currents and correlators are sampled once per driving period.

The chapter is organised in three parts that follow this pipeline: spectral reductions of infinite-time limits [Secs. 2.1 to 2.3](#), model benchmarks that instantiate the resulting diagnostics [Secs. 2.4 to 2.7](#), and the numerical and quantum methods used to evaluate them [Sec. 2.9](#). [Sec. 2.1](#) rewrites infinite-time averages as spectral projections onto degenerate eigenspaces (energies or eigenphases). [Sec. 2.2](#) introduces the spectral densities used throughout (density of states (DOS), local density of states (LDOS), and operator-resolved spectral functions) and links them to equilibrium averages and linear response. [Sec. 2.3](#) then specialises to spin transport, defining current correlators, conductivity, and the Drude weight in both continuous and stroboscopic time. The benchmark part applies this machinery to (i) a non-integrable XYZ chain [Sec. 2.4](#) and (ii) the brickwork XXZ circuit and its disordered extension [Secs. 2.5 and 2.6](#); with models fixed, [Sec. 2.7](#) specifies the lattice current operators and uses a single-impurity setup to contrast ballistic and diffusive channels. Finally, [Sec. 2.9](#) collects the computational ingredients, including polynomially filtered exact diagonalisation method (POLFED) [Sec. 2.9.1](#), symmetry-resolved entanglement [Sec. 2.9.2](#), stochastic trace estimation [Sec. 2.9.3](#), Kernel polynomial method [Sec. 2.9.4](#), quantum moment estimation [Sec. 2.9.5](#), the Quantinuum H1-1 implementation [Sec. 2.9.6](#), and discrete-time spectral reconstructions [Sec. 2.9.7](#).

The material in this chapter that originates from this PhD work is associated with three references. (i) Ref. [1] presents our quantum-classical approach to kernel polynomial method (KPM) (see [Sec. 2.9.5](#)), its application to the XYZ chain to compute the DOS (see [Sec. 2.4](#)), and its implementation on a trapped-ion quantum computer for $N = 18$ spins (see [Sec. 2.9.6](#)); (ii) Ref. [2] covers the discrete-time disordered XXZ circuit and its transport diagnostics (see [Secs. 2.6 and 2.7](#)); and (iii) Ref. [5] (in preparation) develops a method to reconstruct spectral densities of discrete-time correlators (see [Sec. 2.9.7](#)), derives a generic current operator for arbitrary circuit geometries (see [Sec. 2.7.2](#)), and discusses the spectral features of a discrete-time XXZ circuit with an integrability-breaking impurity (see [Sec. 2.7.3](#)). The latter two references are discussed further in [Ch. 3](#).

Throughout this chapter we keep the two settings side-by-side using the dictionary

$$E \leftrightarrow \varphi, \quad \delta(\cdot) \leftrightarrow \delta_{2\pi}(\cdot), \quad \int_{-\infty}^{\infty} dt e^{i\omega t}(\cdot) \leftrightarrow \sum_{t \in \mathbb{Z}} e^{i\omega t}(\cdot), \quad \omega \in [-\pi, \pi), \quad (2.1)$$

where $\delta_{2\pi}(x) = \sum_{m \in \mathbb{Z}} \delta(x - 2\pi m)$ is the 2π -periodic delta distribution.

2.1 Spectral projections

In this section we show that infinite-time averages of observables can be written explicitly using spectral projectors; accordingly, the most natural language for long-time behaviour is the spectrum. For a finite isolated system the evolution is unitary and quasi-periodic, so expectation values need not approach a strict limit. Even so, many observables admit a stable decomposition into (i) a time-independent contribution dictated by exact degeneracies, symmetries, and other conserved structures, and (ii) genuinely time-dependent terms that oscillate and effectively cancel through dephasing when many distinct frequencies are present. Taking an infinite-time average removes the oscillatory sector and retains precisely the stationary component, reducing an infinite-time dynamical problem to a projection (block-diagonalisation) in the eigenbasis. Since our emphasis is on Floquet circuits, we present this construction first for discrete-time evolution generated by a unitary step operator and then state the continuous-time Hamiltonian counterpart.

Infinite-time averages are controlled by spectral projectors. For continuous-time Hamiltonian dynamics the phases are e^{-iEt} with $t \in \mathbb{R}$, while for Floquet dynamics they are $e^{-i\varphi t}$ at stroboscopic times $t \in \mathbb{Z}$. The evolution of an initial state $|\psi(0)\rangle$ is given by

$$\hat{H} |E_i\rangle = E_i |E_i\rangle, \quad |\psi(0)\rangle = \sum_i c_i |E_i\rangle, \quad \Rightarrow \quad |\psi(t)\rangle = \sum_i c_i e^{-iE_i t} |E_i\rangle, \quad t \in \mathbb{R}, \quad (2.2)$$

$$\hat{U} |\varphi_i\rangle = e^{-i\varphi_i} |\varphi_i\rangle, \quad |\psi(0)\rangle = \sum_i c_i |\varphi_i\rangle, \quad \Rightarrow \quad |\psi(t)\rangle = \sum_i c_i e^{-i\varphi_i t} |\varphi_i\rangle, \quad t \in \mathbb{Z}. \quad (2.3)$$

For any observable $\hat{A}$ this gives, respectively,

$$\langle \hat{A}(t) \rangle = \begin{cases} \sum_{i,j} c_i^* c_j e^{it(E_i - E_j)} A_{ij}, & t \in \mathbb{R}, & A_{ij} = \langle E_i | \hat{A} | E_j \rangle, \\ \sum_{i,j} c_i^* c_j e^{it(\varphi_i - \varphi_j)} A_{ij}, & t \in \mathbb{Z}, & A_{ij} = \langle \varphi_i | \hat{A} | \varphi_j \rangle. \end{cases} \quad (2.4)$$

Define Cesàro time averages by

$$\overline{A} = \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T dt \langle \hat{A}(t) \rangle, \quad \overline{A} = \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} \langle \hat{A}(t) \rangle. \quad (2.5)$$

In both cases, the average kills all oscillatory terms and retains only the contributions from degenerate phases: $E_i = E_j$ for $\hat{H}$, and $\varphi_i = \varphi_j$ modulo 2π for $\hat{U}$. Writing the spectral decompositions $\hat{H} = \sum_\alpha E_\alpha \hat{P}_\alpha$ and $\hat{U} = \sum_\alpha e^{-i\varphi_\alpha} \hat{P}_\alpha$, one obtains the same

projector formula

$$\bar{A} = \sum_{\alpha} \hat{P}_{\alpha} \hat{A} \hat{P}_{\alpha}, \quad \bar{A} = \langle \psi(0) | \bar{A} | \psi(0) \rangle. \quad (2.6)$$

If the spectrum is nondegenerate this reduces to the diagonal ensemble: $\bar{A} = \sum_i |c_i|^2 A_{ii}$.

In practice one should resolve exact symmetries before forming $\hat{P}_{\alpha}$ (e.g. work within fixed $U(1)$ charge sectors): otherwise independent symmetry sectors generate spurious degeneracies and therefore spurious stationary contributions. With symmetries resolved, Eq. (2.6) shows that the infinite-time average is obtained by projecting onto degenerate eigenspaces, so infinite-time stationary behaviour is determined by the spectrum (and its exact degeneracies) together with how strongly the observable couples to the corresponding projectors. This is the basic reason why infinite-time equilibration and transport can be phrased as spectral problems (see e.g. [57, 69, 75–77]).

In large systems, the $i \neq j$ terms typically dephase because they add up many oscillations at Bohr frequencies $\omega_{ij} = E_i - E_j$ (or $\omega_{ij} = \varphi_i - \varphi_j$ modulo 2π) whose phases do not remain aligned. Heuristically, they sample a dense, effectively incommensurate set of frequencies. The main exceptions are (i) exact degeneracies or unresolved symmetry sectors, which create many $\omega_{ij} = 0$ contributions, and (ii) singular spectral weight tied to conserved overlaps. For currents (or operators with conserved components), such singular weight appears as a Drude contribution: in the thermodynamic limit (and most cleanly with periodic boundaries) as a $\delta(\omega)$ singularity at $\omega = 0$, while for finite open chains it is redistributed into low-frequency peaks at $\omega \sim 1/N$ (the finite-size manifestation of Mazur-protected ballistic weight; see Part I). The remaining regular low-frequency continuum controls dissipative transport.

A complementary viewpoint is that many dynamical amplitudes are Fourier transforms of spectral measures. For example, in the Hamiltonian case the return amplitude may be written as

$$\langle \psi(0) | \psi(t) \rangle = \begin{cases} \int e^{-iEt} d\mu(E), & t \in \mathbb{R}, & d\mu(E) = \sum_i |c_i|^2 \delta(E - E_i) dE, \\ \int_{-\pi}^{\pi} e^{-it\varphi} d\nu(\varphi), & t \in \mathbb{Z}, & d\nu(\varphi) = \sum_i |c_i|^2 \delta(\varphi - \varphi_i) d\varphi. \end{cases} \quad (2.7)$$

More generally, operator autocorrelators are Fourier transforms of operator-resolved spectral densities.

With the projection identities in hand, the remainder of this section assembles the spectral diagnostics that are used throughout Part I to interpret infinite-time behaviour in both Hamiltonian and Floquet settings. We begin in Sec. 2.1.1 by characterising eigenvalue (or eigenphase) distributions and their local correlations, highlighting unfolding-free measures such as the adjacent gap ratio and stressing that all statistics

must be computed in fully resolved symmetry sectors. In [Sec. 2.1.2](#) we then summarise how eigenvectors enter long-time physics via (Floquet) ETH, viewed as a statement about matrix elements and their frequency dependence, and we indicate how the resulting low-frequency structure controls relaxation and transport. Complementary probes of eigenstate structure (participation and entanglement diagnostics, correlations) are discussed in [Sec. 2.1.3](#). Finally, working directly with spectral densities, the total DOS, LDOS, and operator-resolved spectral functions defined in [Sec. 2.2](#), which provide the practical inputs for the transport diagnostics introduced in [Part I](#) and specialised to currents in [Sec. 2.3](#).

2.1.1 Eigenvalue distribution

The distribution of eigenvalues (or eigenphases) and their short-range correlations fixes the finite-size time and energy scales that govern dephasing, revivals, and slow beats. After resolving all exact symmetries, one typically encounters two canonical cases: Wigner-Dyson statistics, characteristic of ergodic spectra with level repulsion, and Poisson statistics, characteristic of integrable or localised spectra. In the Poisson case, small gaps occur with appreciable probability, so slow beating envelopes can persist to long times at finite size. By contrast, Wigner-Dyson level repulsion suppresses anomalously small gaps, which tends to enhance dephasing and shorten the beating time scales associated with near-degeneracies.

Eigenvalues alone do not determine transport: transport coefficients are controlled by how conserved densities and currents overlap with the slow modes. Level statistics is therefore best viewed as a coarse indicator of ergodicity versus integrability or localisation, and should be complemented by operator-resolved diagnostics (current matrix elements and low-frequency spectral functions), which we implement in [Sec. 2.2](#) and specialise to currents in [Sec. 2.3](#).

A standard unfolding-free diagnostic of local level correlations is the adjacent gap ratio [\[72\]](#). For an ordered set of levels E_i define spacings $\delta E_i = E_{i+1} - E_i$ and

$$r_i = \frac{\min\{\delta E_{i-1}, \delta E_i\}}{\max\{\delta E_{i-1}, \delta E_i\}}. \quad (2.8)$$

For Floquet eigenphases φ_i , spacings must be defined on the unit circle. A practical convention is to choose a cut φ_0 , unwrap phases into $(\varphi_0, \varphi_0 + 2\pi]$, sort them, and compute spacings in that interval (with the wrap-around treated consistently).

In all cases, r must be computed within each exact symmetry block (fixed magnetisation, momentum, parity, etc.), otherwise independent blocks create artificial near-degeneracies and bias $\langle r \rangle$ downward. For uncorrelated (Poissonian) spectra one

finds [72] $\langle r \rangle_{\text{Poisson}} = 2 \ln 2 - 1 \approx 0.386$, while for random-matrix statistics one obtains universal values depending on antiunitary symmetries [78]. In Hamiltonian systems with time-reversal symmetry (Gaussian orthogonal ensemble (GOE), or circular orthogonal ensemble (COE) in the Floquet case), $\langle r \rangle_{\text{GOE/COE}} \simeq 0.5359$, while without time-reversal symmetry (Gaussian unitary ensemble (GUE), or circular unitary ensemble (CUE) for Floquet unitaries), $\langle r \rangle_{\text{GUE/CUE}} \simeq 0.6027$. In the following, we use $\langle r \rangle$ as a compact diagnostic for distinguishing ergodic (circular-ensemble-like) and nonergodic (Poisson-like) regimes of $U(1)$ circuits, and we relate these static signatures to dynamical transport properties. For Floquet-specific ETH/MBL contexts and eigenphase statistics, see e.g. [79].

2.1.2 Eigenstate thermalisation and long-time transport

Eigenvectors provide the second pillar of the spectral viewpoint: in ergodic many-body systems, few-body observables in individual eigenstates look thermal, as captured by the eigenstate thermalisation hypothesis (ETH) [46, 59, 73, 80, 81]. A common formulation is the ETH ansatz for matrix elements of a local operator $\hat{A}$ in the Hamiltonian eigenbasis,

$$A_{ij} = A(\bar{E}) \delta_{ij} + e^{-s(\bar{E})/2} f_A(\bar{E}, \omega) R_{ij}, \quad \bar{E} = \frac{E_i + E_j}{2}, \quad \omega = E_j - E_i, \quad (2.9)$$

where $s(\bar{E}) = \log g(\bar{E})$ is the microcanonical entropy¹, f_A is a smooth envelope, and R_{ij} are zero-mean, unit-variance random variables. We indicate the entropy by the lowercase s to distinguish it from the spectral functions defined in Sec. 2.2.

The smooth function $A(E)$ is identified with the microcanonical expectation value

$$A(E) = \langle \hat{A} \rangle_{\text{mc}}(E) = \frac{\text{Tr}[\hat{A} \hat{P}_{E,\Delta}]}{\text{Tr}[\hat{P}_{E,\Delta}]}, \quad \hat{P}_{E,\Delta} = \sum_{i: |E_i - E| \leq \Delta} |E_i\rangle \langle E_i|, \quad (2.10)$$

so $A_{ii} \approx A(E_i)$ varies smoothly with energy, while off-diagonal elements are exponentially small in system size. If the initial state has support in a narrow energy window, the diagonal ensemble implies $\bar{A} \approx A(E)$, i.e. few-body observables relax to their equilibrium values for broad classes of initial states [46, 59].

In Floquet systems without energy conservation, the corresponding notion of “thermal” is typically infinite temperature within the symmetry sector. If $\hat{Q}$ is a conserved charge

¹Where $g(E)$ is the density of states (DOS), see Sec. 2.2.

with sector projector $\hat{P}_q$, the infinite-temperature expectation in that sector is

$$A_\infty(q) = \frac{\text{Tr}[\hat{A} \hat{P}_q]}{\text{Tr}[\hat{P}_q]}, \quad (2.11)$$

and Floquet-ETH asserts that diagonal matrix elements in the eigenbasis of $\hat{U}$ concentrate near $A_\infty(q)$, while off-diagonal elements have a smooth envelope as a function of the eigenphase difference

$$\omega = \text{wrap}_{(-\pi, \pi]}(\varphi_j - \varphi_i). \quad (2.12)$$

The off-diagonal structure is not a technical detail: for currents and conserved densities it governs the decay of correlators and the low-frequency spectral weight entering Kubo-type formulas, while conservation laws constrain the $\omega \rightarrow 0$ hydrodynamic regime.

It is useful to distinguish strong versus weak ETH for the diagonal sector. Fix an energy window $\mathcal{W}_{E,\Delta} = \{i : E_i \in [E - \Delta, E + \Delta]\}$ and define $\delta_A(i) = |A_{ii} - A(E_i)|$. Strong ETH asserts $\max_{i \in \mathcal{W}_{E,\Delta}} \delta_A(i) \rightarrow 0$ as $N \rightarrow \infty$ [81, 82], while weak ETH allows rare outliers but requires their fraction to vanish. Integrable and many-body localised systems violate ETH in characteristic ways [46, 58, 83], which is why combining eigenstate probes with operator-resolved spectral functions is diagnostically powerful.

2.1.3 Structure of eigenstates

The structure of eigenstates determines how a local operator or an initial state decomposes across the spectrum. If a local excitation has appreciable weight on an extensive set of eigenstates, interference among many incommensurate frequencies promotes decay and spreading; if it overlaps only a few eigenstates, dynamics remains constrained and transport is suppressed.

Hilbert-space measures such as inverse participation ratios and participation numbers quantify how many basis states significantly contribute to a given eigenvector, but they are basis-dependent; throughout we mean the computational (product) basis when using such measures. Entanglement-based diagnostics are basis-independent and in fact closely connected to ETH. For a bipartition into regions A and B with $N_A + N_B = N$, define $\hat{\rho}_A^{(i)} = \text{Tr}_B[|E_i\rangle\langle E_i|]$ and the von Neumann entropy; for Hilbert spaces of size $D_A = 2^{N_A}$ and $D_B = 2^{N_B}$.

$$s^{(i)}(A) = -\text{Tr}[\hat{\rho}_A^{(i)} \log \hat{\rho}_A^{(i)}]. \quad (2.13)$$

In Floquet circuits the same eigenstate diagnostics apply to eigenstates of $\hat{U}$. In particular, for a Floquet eigenstate $|\varphi_i\rangle$ we define the half-system von Neumann

entropy

$$s^{(i)}(A) = -\text{Tr}[\hat{\rho}_A^{(i)} \ln(\hat{\rho}_A^{(i)})], \quad (2.14)$$

where $\hat{\rho}_A^{(i)} = \text{Tr}_B[|\varphi_i\rangle\langle\varphi_i|]$ is the reduced density matrix for subsystem A . Averaging s_i over eigenstates in a fixed magnetisation sector yields the mean eigenstate entanglement entropy $\langle s \rangle$. In ergodic phases satisfying ETH, highly excited eigenstates exhibit a volume law,

$$s^{(i)}(A) \approx s(e_i) N_A, \quad (2.15)$$

and approach the random-state benchmark appropriate to the Hilbert-space dimension under the imposed symmetries. In fixed magnetisation sectors, the relevant benchmark is the symmetry-resolved Page value, obtained by replacing D_A, D_B by the corresponding sector dimensions.

In the absence of symmetry constraints a random pure state on $\mathcal{H}_A \otimes \mathcal{H}_B$ with dimensions $D_A \leq D_B$ has (unconstrained) Page value [84]

$$s_{\text{Page}}^{(\text{full})} \approx \log D_A - \frac{D_A}{2D_B}, \quad (2.16)$$

which for an equal bipartition ($D_A = D_B = 2^{N/2}$) yields $s_{\text{Page}}^{(\text{full})} \approx \frac{N}{2} \log 2 - \frac{1}{2}$. Instead, when numerics are performed in a fixed total magnetisation sector M (as in Sec. 2.6), the natural random-state benchmark is the Page value restricted to that charge sector. The exact finite-size expression has been derived in Ref. [85]; here N_A is the subsystem size with $N_A + N_B = N$, and M_A is the magnetisation in A ,

$$s_{\text{Page}}^{(M)} = \frac{D_{A,M_A} D_{B,M_B}}{D_M} \left(\sum_{M_A} \bar{s}_{M_A} - \sum_{M_A} (\Psi(D_{A,M_A} + 1) - \Psi(D_M + 1)) \right), \quad (2.17)$$

with Ψ the digamma function, with $D_{A,M_A} = \dim \mathcal{H}_{A,M_A}$ and $D_M = \sum_{M_A} D_{A,M_A} D_{B,M_B}$, and where $\bar{s}_{M_A}$ is the ordinary Page entropy from Eq. (2.16). When assessing Page-like saturation in the ergodic regime we therefore normalise the mean eigenstate entropies by this fixed-magnetisation Page value, $\langle s \rangle / s_{\text{Page}}$. In many-body localised phases, eigenstates violate ETH and exhibit an area law (up to known logarithmic corrections in some settings), $s^{(i)}(A) = \mathcal{O}(1)$, independent of N_A in the thermodynamic limit.

A related correlation diagnostic is the quantum mutual information between disjoint regions A and B ,

$$I(A : B) = s(\hat{\rho}_A) + s(\hat{\rho}_B) - s(\hat{\rho}_{AB}), \quad (2.18)$$

where $\hat{\rho}_A, \hat{\rho}_B, \hat{\rho}_{AB}$ are reduced density matrices of the eigenstate under consideration. In ergodic phases, $I(A : B)$ decays rapidly with separation, while in MBL phases it reveals residual correlations mediated by quasi-local integrals of motion.

2.2 Spectral densities

Infinite-time limits are controlled by spectral projections (Sec. 2.1): time averaging removes oscillatory components and retains only the stationary contributions selected by degeneracies and symmetries. For transport, one must additionally resolve *how much* spectral weight a given state or operator places on different frequencies, especially near $\omega \simeq 0$.

We therefore work with *spectral densities*: the global density of states (DOS) g , its state- and operator-resolved counterparts g_ψ and g_A (often grouped under LDOS), the linear density h_A relevant for one-point averages, and dynamical (two-point) spectral densities $S_A(\omega)$ obtained as Fourier transforms of equilibrium autocorrelators $\langle \hat{A}(t) \hat{A}^\dagger \rangle$. These dynamical spectral densities are the natural inputs for linear-response transport and the Drude-weight diagnostics developed below. At finite size these are sums of δ -peaks; we treat them as distributional densities with respect to dE (Hamiltonian) or $d\varphi$ on the unit circle (Floquet), using g as generic notation for the positive densities $\{g, g_\psi, g_A\}$.

Computationally, full exact diagonalisation (ED) is quickly out of reach ($\mathcal{O}(D^3)$ time and $\mathcal{O}(D^2)$ memory with $D = 2^N$). For most long-time and transport questions this is unnecessary: the relevant information enters through these weighted spectral densities, which can be reconstructed from moments (Chebyshev/KPM for Hamiltonians and Fourier moments for Floquet unitaries; see Sec. 2.9.4).

Let $\hat{H} |E_i\rangle = E_i |E_i\rangle$ on a Hilbert space of size D (for N spins, $D = 2^N$). The density of states (DOS) is the normalised spectral measure

$$g(E) = \frac{1}{D} \sum_i \delta(E - E_i) = \frac{1}{D} \text{Tr}[\delta(E - \hat{H})], \quad \int dE g(E) = 1. \quad (2.19)$$

The Floquet analogue replaces energies by eigenphases. Writing $\hat{U} |\varphi_i\rangle = e^{-i\varphi_i} |\varphi_i\rangle$, $\varphi_i \in [-\pi, \pi)$, the eigenphase DOS is

$$g(\varphi) = \frac{1}{D} \sum_i \delta_{2\pi}(\varphi - \varphi_i), \quad \int_{-\pi}^{\pi} d\varphi g(\varphi) = 1. \quad (2.20)$$

Both objects quantify how densely the spectrum fills its natural domain (line vs circle), but they contain no information about how spectral weight is distributed across states or operators.

Spatial or state selectivity is introduced by the local density of states (LDOS). In the many-body literature, two related objects are both commonly called LDOS [46]. First, the state-resolved spectral density (energy-resolved weight of a chosen initial state $|\psi\rangle$)

is

$$g_\psi(E) = \sum_i |\langle E_i | \psi \rangle|^2 \delta(E - E_i) = \langle \psi | \delta(E - \hat{H}) | \psi \rangle. \quad (2.21)$$

The Floquet analogue is the eigenphase spectral density of $|\psi\rangle$,

$$g_\psi(\varphi) = \sum_i |\langle \varphi_i | \psi \rangle|^2 \delta_{2\pi}(\varphi - \varphi_i). \quad (2.22)$$

Second, an operator-resolved spectral density weights the spectrum by a chosen operator $\hat{A}$,

$$g_A(E) = \frac{1}{\text{Tr}[\hat{A}^\dagger \hat{A}]} \text{Tr}[\hat{A}^\dagger \delta(E - \hat{H}) \hat{A}], \quad g_A(\varphi) = \frac{1}{\text{Tr}[\hat{A}^\dagger \hat{A}]} \text{Tr}[\hat{A}^\dagger \delta_{2\pi}(\varphi - \hat{\varphi}) \hat{A}], \quad (2.23)$$

where $\hat{\varphi} = \sum_i \varphi_i |\varphi_i\rangle\langle\varphi_i|$ denotes the spectral phase operator on the chosen principal branch. In the energy basis, $g_A(E)$ weights each shell by the total matrix-element weight $\sum_j |A_{ij}|^2$ and is therefore the natural input for quadratic diagnostics (autocorrelators, spectral functions, transport kernels).

For one-point equilibrium averages of $\hat{A}$ itself, one instead needs the *linear* (diagonal) spectral density

$$h_A(E) = \frac{1}{D} \text{Tr}[\hat{A} \delta(E - \hat{H})] = \frac{1}{D} \sum_i A_{ii} \delta(E - E_i), \quad (D = 2^N), \quad (2.24)$$

and, in Floquet systems,

$$h_A(\varphi) = \frac{1}{D} \text{Tr}[\hat{A} \delta_{2\pi}(\varphi - \hat{\varphi})] = \frac{1}{D} \sum_i A_{ii} \delta_{2\pi}(\varphi - \varphi_i). \quad (2.25)$$

Unlike g_A , the functions h_A are not positive and satisfy $\int dE h_A(E) = D^{-1} \text{Tr}[\hat{A}]$ (typically zero for local traceless probes), but they are precisely the objects whose Laplace transform yields $\text{Tr}[\hat{A} e^{-\beta \hat{H}}]$. We compute the DOS and this quantity for an XYZ model (both continuous and discrete time) in [Sec. 2.4](#).

More general spectral functions arise from two-point correlators. For continuous time, starting from the autocorrelator of $\hat{A}$,

$$C_A(t) = C_{AA^\dagger}(t) = \langle \hat{A}(t) \hat{A}^\dagger \rangle = \text{Tr}[\hat{\rho} \hat{A}(t) \hat{A}^\dagger], \quad \hat{A}(t) = e^{i\hat{H}t} \hat{A} e^{-i\hat{H}t}, \quad (2.26)$$

its Fourier transform defines the corresponding spectral density. Evaluating the trace

in the energy (or eigenphase) eigenbasis yields the standard Lehmann representation,

$$S_A(\omega) = S_{AA^\dagger}(\omega) = \frac{1}{2\pi} \begin{cases} \int_{-\infty}^{\infty} dt e^{i\omega t} C_A(t) = \sum_{i,j} p_i |\langle E_j | \hat{A} | E_i \rangle|^2 \\ \quad \times \delta(\omega - (E_j - E_i)), \quad \omega \in \mathbb{R}, \\ \sum_{t \in \mathbb{Z}} e^{i\omega t} C_A(t) = \sum_{i,j} p_i |\langle \varphi_j | \hat{A} | \varphi_i \rangle|^2 \\ \quad \times \delta_{2\pi}(\omega - (\varphi_j - \varphi_i)), \quad \omega \in [-\pi, \pi), \end{cases} \quad (2.27)$$

respectively for Hamiltonian and Floquet dynamics, where $\hat{\rho} = \sum_i p_i |E_i\rangle\langle E_i|$ (or $\hat{\rho} = \sum_i p_i |\varphi_i\rangle\langle \varphi_i|$ in the Floquet case). In both settings, choosing $\hat{A}$ to be a density or current yields dynamical structure factors and current spectral densities (see the following [Sec. 2.3](#) for the latter).

For a Gibbs state $\hat{\rho} = \hat{\rho}_\beta \propto e^{-\beta \hat{H}}$, the KMS condition implies detailed balance,

$$S_A(-\omega) = e^{-\beta\omega} S_{A^\dagger}(\omega). \quad (2.28)$$

In particular, if $\hat{A} = \hat{A}^\dagger$ then $S_A(-\omega) = e^{-\beta\omega} S_A(\omega)$. In linear response one instead uses the retarded correlator $\chi_J^R(t) = -i\Theta(t)\langle[\hat{J}(t), \hat{J}(0)]\rangle$; its imaginary part is related to S_J by the fluctuation-dissipation theorem (with the Fourier conventions fixed above). We state explicitly, in each later use, whether we work with S_J (spectral density of the current correlator) or with χ_J^R (retarded response); the long-time Drude channel extracted below can be expressed in either language.

2.3 Currents and Drude weight

With the framework of [Secs. 2.1](#) and [2.2](#) in hand, we turn to transport. Transport is obtained by specialising the spectral functions of [Sec. 2.2](#) to a conserved quantity and its associated current, so that ballistic channels (Drude weight) and long-time tails can be read directly from the infinite-time behaviour of current autocorrelators and their zero-frequency structure. In Hamiltonian chains and in U(1)-symmetric Floquet circuits the common conserved charge is the magnetisation $\hat{S}^z = \sum_n \hat{s}_n^z$; as mentioned already we focus solely on spin transport. For a concrete illustration of why spectral-chaos diagnostics and long-time transport can decouple at finite size, see the single-impurity benchmark in [Sec. 2.7](#).

Conservation of magnetisation fixes the notion of current through a lattice continuity equation. Let $\hat{s}_n^z(t)$ denote the Heisenberg-evolved local charge density, either under

$\hat{H}$ or under $\hat{U}$. For Hamiltonian dynamics one has

$$\frac{d}{dt} \hat{s}_n^z(t) = i[\hat{H}, \hat{s}_n^z(t)] = \hat{j}_{n-1}^z(t) - \hat{j}_n^z(t), \quad (2.29)$$

where $\hat{j}_n^z$ is a local current. For Floquet dynamics, charge conservation is stroboscopic and the continuity equation reads

$$\frac{\Delta}{\Delta t} \hat{s}_n^z(t) = \hat{U}^\dagger \hat{s}_n^z(t) \hat{U} - \hat{s}_n^z(t) = \hat{j}_{n-1}^z(t) - \hat{j}_n^z(t), \quad (2.30)$$

which yields a natural per-cycle current compatible with the circuit architecture. In both settings the total current is $\hat{J}^z = \sum_n \hat{j}_n^z$.

We diagnose transport via the equilibrium autocorrelator

$$C_J(t) = \frac{1}{N} \langle \hat{J}^z(t) \hat{J}^z(0) \rangle, \quad (2.31)$$

and its frequency-domain spectral density $S_J(\omega)$. For Hamiltonian dynamics one uses a Fourier transform on $\mathbb{R}$, while for Floquet dynamics one uses the discrete-time Fourier transform on the Brillouin zone $\omega \in [-\pi, \pi)$:

$$S_J(\omega) = \frac{1}{2\pi} \begin{cases} \int_{-\infty}^{\infty} dt e^{i\omega t} C_J(t), & \text{Hamiltonian (continuous time),} \\ \sum_{t \in \mathbb{Z}} e^{i\omega t} C_J(t), & \text{Floquet (stroboscopic time).} \end{cases} \quad (2.32)$$

Both are special cases of the Lehmann representations introduced earlier, with $\hat{A} = \hat{J}^z$.

At finite temperature, linear response defines the (spin) conductivity $\sigma(\omega)$ from the retarded current correlator [43, 50, 70],

$$\chi_J^R(t) = -\frac{i}{N} \theta(t) \langle [\hat{J}^z(t), \hat{J}^z(0)] \rangle, \quad \sigma(\omega) = \frac{1}{i\omega} \chi_J^R(\omega), \quad (2.33)$$

with $\chi_J^R(\omega) = \int_{-\infty}^{\infty} dt e^{i\omega t} \chi_J^R(t)$ for Hamiltonians, and $\chi_J^R(\omega) = \sum_{t \in \mathbb{Z}} e^{i\omega t} \chi_J^R(t)$ for Floquet stroboscopic time. For a Gibbs state $\hat{\rho}_\beta \propto e^{-\beta \hat{H}}$, the Kubo-Martin-Schwinger (KMS) condition yields the fluctuation-dissipation relation

$$\text{Re } \sigma_{\text{reg}}(\omega) = \frac{\pi}{\omega} (1 - e^{-\beta\omega}) S_J(\omega), \quad (\omega \neq 0), \quad (2.34)$$

so $\sigma_{\text{reg}}(\omega)$ is the absolutely continuous part of $\text{Re } \sigma(\omega)$ away from $\omega = 0$.

The infinite-time (ballistic) channel is the non-decaying component of $C_J(t)$. A convenient finite-size diagnostic is the Cesàro average of Eq. (2.5) specialised to

currents:

$$C_\infty = \lim_{T \rightarrow \infty} \frac{1}{T} \begin{cases} \int_0^T dt C_J(t), & \text{Hamiltonian,} \\ \sum_{t=0}^{T-1} C_J(t), & \text{Floquet.} \end{cases} \quad (2.35)$$

In frequency space this corresponds to a zero-frequency singular contribution [43, 59]. In the Hamiltonian case one writes

$$\text{Re } \sigma(\omega) = 2\pi D \delta(\omega) + \sigma_{\text{reg}}(\omega), \quad D = \frac{\beta}{2} C_\infty, \quad (2.36)$$

where the second equality follows from Eq. (2.34) and the fact that C_∞ produces a $\delta(\omega)$ component in $S_J(\omega)$. For Floquet circuits the analogue is a $\delta_{2\pi}(\omega)$ contribution at $\omega = 0$, reflecting the 2π -periodicity of the stroboscopic spectrum; the finite- β form Eq. (2.34) applies when the stationary state is Gibbs with respect to a (prethermal) effective Hamiltonian, while in generic driven settings we use $S_J(\omega)$ and C_∞ directly as the primary diagnostics.

This finite-size phenomenology is precisely the content of Eq. (2.36). In the thermodynamic limit, ballistic transport corresponds to a nonzero Drude weight D , i.e. to a $\delta(\omega)$ contribution in $\text{Re } \sigma(\omega)$. At finite N (especially for open boundaries) the $\delta(\omega)$ term is not seen as an exact singularity; rather, its weight is distributed among low-frequency peaks whose characteristic positions scale as $\omega \sim 1/N$ (as expected from Mazur-type bounds) and therefore accumulate towards $\omega \rightarrow 0$ only as $N \rightarrow \infty$. Consequently, the robust quantity to track is the *integrated* low-frequency weight, the finite- N proxy for $2\pi D$ in Eq. (2.36), not the peak height or width. The physical mechanism behind $C_\infty > 0$ is the existence of nontrivial conserved or quasi-conserved operators overlapping with the current. In an eigenbasis the same mechanism is transparent. For Hamiltonians, only matrix elements of $\hat{J}^z$ within degenerate energy subspaces survive the $T \rightarrow \infty$ averaging. For Floquet circuits, the surviving contributions come from degenerate eigenphases modulo 2π . Thus $C_\infty \neq 0$ signals a nonzero overlap of $\hat{J}^z$ with conserved (or quasi-conserved) operators, preventing complete decay of current correlations. We return to this point in Sec. 2.7.

2.4 Spectral densities of a XYZ chain

We now apply the spectral densities of Sec. 2.2 to a concrete many-body model, chosen to be generic (non-integrable) while remaining simple enough to serve as a numerical benchmark. The objective of this section is methodological. We demonstrate, side-by-side for Hamiltonian and Floquet dynamics, how to reconstruct global and operator-resolved spectral data, and how to use these reconstructions as inputs to infinite-




Figure 2.1: Density of states of the non-integrable XYZ chain Eq. (2.37), computed by ED (histogram) and by a moment-based reconstruction (yellow line), together with the proof-of-principle quantum KPM estimates discussed in Sec. 2.9.6 (blue line), for $N = 12$ (a) and $N = 18$ (b).

time diagnostics and equilibrium averages. In addition to classical reconstructions, we will also use this same model as a proof-of-principle benchmark for a quantum implementation of moment estimation on Quantinuum’s H1-1 device of Ref. [1]; the corresponding hardware data are discussed in Sec. 2.9.6 and included in the DOS comparisons below.

2.4.1 Continuous time

We consider a spin- $\frac{1}{2}$ XYZ Heisenberg chain with a staggered longitudinal interaction,

$$\hat{H} = \sum_{n=1}^N J_x \hat{s}_n^x \hat{s}_{n+1}^x + J_y \hat{s}_n^y \hat{s}_{n+1}^y + (J_z + (-1)^n \Lambda) \hat{s}_n^z \hat{s}_{n+1}^z, \quad (2.37)$$

with PBC and parameters $J_x = 1$, $J_y = 1/3$, $J_z = \Lambda = 1/2$. The bond-alternating zz coupling breaks the integrable structure of the uniform XYZ chain and yields generic spectral properties (beyond the exact symmetries discussed in Sec. 2.1), as explored for related staggered models in Ref. [86]. We compute the DOS $g(E)$ as defined in Eq. (2.19), and compare a direct histogram from ED with moment-based reconstructions in Fig. 2.1 for $N = 12$ and $N = 18$. The yellow curves show the classical KPM results, while the blue curves correspond to data obtained from the quantum device. The point is not the specific reconstruction algorithm, but the representation: for long-time questions it is often sufficient to reconstruct spectral densities without resolving individual eigenpairs.

For Hamiltonians that are sums of many bounded, finite-range terms, the global DOS typically develops a smooth, approximately Gaussian envelope at large N [87]. In practice, low-order moments already fix the bulk envelope, while increasingly high




Figure 2.2: (a) DOS reconstruction and (b) an operator-resolved spectral density for $\hat{A} = \hat{s}_1^z \hat{s}_2^z$, comparing ED with moment-based reconstructions, system size $N = 12$. (c) Canonical average $\langle \hat{s}_1^z \hat{s}_2^z \rangle_\beta$ obtained from the spectral densities via Eq. (2.38).

orders are mainly required to sharpen fine spectral features.

This dichotomy is also visible in the quantum experiment: on H1-1, we can reliably resolve only a finite band of low-order Chebyshev moments, and the smallest moments beyond a device-dependent cutoff are statistically indistinguishable from zero at fixed shot budget. Nevertheless, for $N = 12$ we find that moments up to $m \leq 7$ already suffice to reconstruct the bulk DOS at the resolution used in Fig. 2.1, while for $N = 18$ the accessible range extends to approximately $m \leq 11$ on the emulator and only partially on hardware (details in Sec. 2.9.6). This is precisely the regime in which moment methods are useful: one trades spectral fine structure for robustness and scalability.

As discussed in Sec. 2.2, to resolve the spectrum with respect to a chosen observable, one can use operator-resolved spectral densities like $g_A(E)$ of Eq. (2.23) and $h_A(E)$ of Eq. (2.24). Canonical averages can be written purely as energy integrals,

$$\langle \hat{A} \rangle_\beta = \frac{\text{Tr}[\hat{A} e^{-\beta \hat{H}}]}{\text{Tr}[e^{-\beta \hat{H}}]} = \frac{\int dE e^{-\beta E} h_A(E)}{\int dE e^{-\beta E} g(E)}, \quad (2.38)$$

with $g(E)$ and $h_A(E)$ as in Eqs. (2.19) and (2.24), respectively. In Fig. 2.2 we show a representative reconstruction for $N = 12$ with $\hat{A} = \hat{s}_1^z \hat{s}_2^z$, including the resulting canonical average obtained from Eq. (2.38); technical implementation details are deferred to the methods chapters.

2.4.2 Discrete time

A discrete-time analogue of the staggered XYZ chain is obtained by exponentiating each bond term of Eq. (2.37) into a two-site gate $\hat{u}_{n,n+1}(\tau) = e^{-i\tau \hat{h}_{n,n+1}}$ and assembling a two-layer brickwork step,

$$\hat{U}(\tau) = \prod_{n \text{ odd}} \hat{u}_{n,n+1}(\tau) \prod_{n \text{ even}} \hat{u}_{n,n+1}(\tau), \quad (2.39)$$




Figure 2.3: (a) Eigenphase DOS $g(\varphi)$ of the staggered XYZ brickwork Floquet circuit Eq. (2.39). Bars: ED histogram of Floquet eigenphases $\{\varphi_i\}$. Line: Fourier reconstruction from trace moments $s_m = \text{Tr}[\hat{U}^m]/D$, damped by a circular heat kernel and truncated at order M . (b) Eigenphase-resolved finite-temperature spectral density $\tilde{S}_A^{(\beta)}(\varphi)$ on the unit circle for a representative local two-site observable $\hat{A}$ (here chosen as a non-diagonal bond operator). Bars: ED estimate. Line: heat-kernel filtered Fourier reconstruction from thermal autocorrelator moments $s_{A,m}^{(\beta)}$, estimated by thermal typicality with R random probes (parameters in Sec. 2.9.7).

with PBC. The spectrum of $\hat{U}$ consists of eigenphases $\varphi_i \in [-\pi, \pi)$ defined by $\hat{U}|\varphi_i\rangle = e^{-i\varphi_i}|\varphi_i\rangle$, and the corresponding global eigenphase DOS is the normalised counting measure on these points. In Fig. 2.3(a) we compare the ED histogram of $\{\varphi_i\}$ to a moment reconstruction obtained from the normalised traces $s_m = \text{Tr}[\hat{U}^m]/D$ (with $D = 2^N$): truncating the Fourier series at order M produces Gibbs-type ringing on the circle, which we suppress by multiplying each harmonic by a heat-kernel factor $\gamma_m = \exp[-(m\sigma)^2/2]$. The parameters (M, σ) therefore set an effective angular resolution, trading fine structure for a stable, smooth spectral density.

For the reconstruction parameters (cutoff M , heat-kernel width σ , grids and histogram binning) used in Fig. 2.3, see Sec. 2.9.7.

Panel (b) of Fig. 2.3 illustrates an operator-resolved spectral object that is specific to discrete time: the eigenphase-difference density governing a Floquet autocorrelation function. For the observable $\hat{A} = \hat{\sigma}_1^x \hat{\sigma}_2^y - \hat{\sigma}_1^y \hat{\sigma}_2^x$ and inverse temperature β (defined with respect to the reference Hamiltonian $\hat{H}$ of Eq. (2.37)), we form the thermal Floquet moments

$$s_{A,m}^{(\beta)} = \frac{\text{Tr}[e^{-\beta\hat{H}} \hat{U}^{-m} \hat{A} \hat{U}^m \hat{A}]}{\text{Tr}[e^{-\beta\hat{H}}]}, \quad m \in \mathbb{Z}, \quad (2.40)$$

and reconstruct their (heat-kernel smoothed) Fourier transform on the unit circle,

$$S_A^{(\beta)}(\varphi) \approx \frac{1}{2\pi} \left[s_{A,0}^{(\beta)} + 2 \sum_{m=1}^M \gamma_m \text{Re}(s_{A,m}^{(\beta)} e^{im\varphi}) \right], \quad \tilde{S}_A^{(\beta)}(\varphi) = \frac{S_A^{(\beta)}(\varphi)}{s_{A,0}^{(\beta)}}. \quad (2.41)$$

This object is a (heat-kernel smoothed) *dynamical spectral density* on the unit circle: it is

defined so that its Fourier coefficients reproduce the thermal stroboscopic autocorrelator moments,

$$s_{A,m}^{(\beta)} = \langle \hat{A}(m) \hat{A} \rangle_\beta = \int_{-\pi}^{\pi} d\varphi S_A^{(\beta)}(\varphi) e^{-im\varphi}, \quad m \in \mathbb{Z}. \quad (2.42)$$

The normalised density $\tilde{S}_A^{(\beta)}(\varphi)$ integrates to unity and resolves how strongly $\hat{A}$ couples to phase differences $\varphi_\alpha - \varphi_\beta$ between Floquet eigenstates, with matrix-element weights and thermal weighting implicit in Eq. (2.40). At the system sizes of Fig. 2.3, these moments can be computed exactly by ED; more generally they can be estimated without diagonalisation by thermal typicality, which is the route used for the smooth curve in Fig. 2.3(b). Together, panels (a) and (b) exemplify the two complementary roles of spectral densities in discrete time: a global eigenphase density (the eigenphase DOS) set by $\text{Tr}[\hat{U}^m]$, and an observable-resolved spectral measure built from thermal Floquet correlators.

The reconstruction scheme used in Sec. 2.9.7 makes explicit a step that is often left implicit in Floquet spectral analysis: regularisation. To our knowledge, this moment-based formulation for eigenphase-resolved spectra has not been stated in this form in the existing literature, so we summarise the underlying idea and its relation to standard practice.

In experiments and numerics, eigenphase-resolved spectra are commonly obtained from a finite stroboscopic time series by a discrete Fourier transform, sometimes with a window to reduce leakage and ringing (e.g. the cosine window in Ref. [88], and related reconstructions in Refs. [89, 90]). In that approach, smoothing is encoded in the truncation/window: the output is the true 2π -periodic spectral density convolved with a finite-time kernel (Dirichlet for hard truncation, or a window-dependent modification).

Here we instead treat the stroboscopic correlators as Fourier moments $s_{A,m}^{(\beta)}$ of a 2π -periodic spectral measure and reconstruct $S_A^{(\beta)}(\varphi)$ from a damped Fourier series,

$$S_A^{(\beta)}(\varphi) \approx \sum_{m=-M}^M \gamma_{|m|} s_{A,m}^{(\beta)} e^{-im\varphi}, \quad (2.43)$$

with a prescribed damping profile γ_m . This makes the estimator well-defined at finite M and separates modelling choices (the kernel) from the data (the moments). Choosing $\gamma_m = \exp[-(m\sigma)^2/2]$ corresponds to convolution with the circular heat kernel on S^1 . The single parameter σ then sets the eigenphase resolution and suppresses high- m components, which is essential when the moments are noisy (for instance when $s_{A,m}^{(\beta)}$ are obtained stochastically via dynamical typicality; Sec. 2.9.3). Conceptually, this is the unit-circle analogue of kernel damping in Chebyshev KPM: the kernel is part of the estimator and controls the bias-variance tradeoff, while respecting the 2π periodicity of Floquet spectra.

Related ideas appear in the Floquet-heating literature, where one introduces an explicit broadening to define a smooth eigenphase spectral function. For example, Ref. [91] defines $\mathcal{S}(\omega)$ from damped discrete Fourier transforms of stroboscopic connected correlators, using exponential (Lorentzian) or Gaussian (Jacobi-theta) filters. The emphasis here is different: we frame eigenphase reconstruction as a moment expansion on the circle and use the heat kernel as a natural default, giving a translation-invariant smoothing of the periodic $\delta_{2\pi}$ peaks and interfacing directly with moment-estimation primitives for circuit dynamics.

2.5 Integrable discrete-time XXZ chain

As a proof of principle in Sec. 2.4 we adopted a non-integrable staggered XYZ chain, chosen precisely to break integrability. A simpler and paradigmatic benchmark is the clean XXZ spin- $\frac{1}{2}$ chain, which conserves U(1) charge and is integrable by Bethe ansatz. In Hamiltonian form,

$$\hat{H}_{\text{XXZ}} = \sum_{n=1}^N \hat{h}_{n,n+1}, \quad \hat{S}^z = \sum_{n=1}^N \hat{s}_n^z, \quad (2.44)$$

with PBC and where $[\hat{H}_{\text{XXZ}}, \hat{S}^z] = 0$. We parameterise the local interaction as

$$\begin{aligned} \hat{h}_{n,n+1} &= \frac{J}{2} \left(\hat{s}_n^+ \hat{s}_{n+1}^- + \hat{s}_n^- \hat{s}_{n+1}^+ \right) + J_z \hat{s}_n^z \hat{s}_{n+1}^z \\ &= J \left(\hat{s}_n^x \hat{s}_{n+1}^x + \hat{s}_n^y \hat{s}_{n+1}^y \right) + J_z \hat{s}_n^z \hat{s}_{n+1}^z \end{aligned} \quad (2.45)$$

This choice matches the J -convention used later in the disordered circuit (Eq. (2.55) at $\phi_n = h_n = h'_n = 0$), and makes the $J = \pi$ point coincide with the SWAP-type (dual-unitary) limit discussed in Sec. 2.6.1.

A generic brickwork discretisation is obtained by a Suzuki-Trotter (ST) step:

$$\begin{aligned} \hat{U} &= e^{-i\hat{H}_{\text{odd}}} e^{-i\hat{H}_{\text{even}}} \\ &= \prod_{n \text{ odd}} e^{-i\hat{h}_{n,n+1}} \prod_{n \text{ even}} e^{-i\hat{h}_{n,n+1}}. \end{aligned} \quad (2.46)$$

For small step size this approximates the continuous-time XXZ dynamics, but at finite




Figure 2.4: Phase diagram of the integrable discrete-time (brickwork) XXZ circuit. The line $J = \pi$ is dual-unitary (and, within the present U(1)-symmetric gate family, reduces to a generalised SWAP). The diagonal $J = J_z$ is the isotropic (XXX) point. In the easy-plane region $J > J_z$ spin transport is ballistic (nonzero Drude weight), while in the easy-axis region $J < J_z$ it is normal diffusive at half filling; at $J = J_z$ spin dynamics is superdiffusive with Kardar-Parisi-Zhang (KPZ) scaling.

step size it is, in general, *not* guaranteed to inherit algebraic integrability.

The integrable *Floquet* XXZ circuit is instead defined by choosing the two-site gate to be (the braid form of) the six-vertex R -matrix, which satisfies the Yang-Baxter equation. A standard unitary parametrisation is

$$\hat{u}(\lambda, \eta) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & a(\lambda, \eta) & b(\lambda, \eta) & 0 \\ 0 & b(\lambda, \eta) & a(\lambda, \eta) & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}. \quad (2.47)$$

with

$$a(\lambda, \eta) = \frac{\sin \lambda}{\sin(\lambda + i\eta)}, \quad b(\lambda, \eta) = \frac{\sin(i\eta)}{\sin(\lambda + i\eta)}. \quad (2.48)$$

The Floquet propagator is then the brickwork product of these R -matrix gates, $\hat{U}_{\text{XXZ}} = \prod_{\text{odd}} \hat{u} \prod_{\text{even}} \hat{u}$, and inherits an extensive family of local and quasi-local conserved charges generated by commuting transfer matrices. We make this transfer-matrix construction explicit and use the first non-trivial local charge as a probe of impurity-induced prethermal decay in [Sec. 3.7](#). In the small-step limit ($\lambda \rightarrow 0$ with η fixed) the corresponding stroboscopic generator reduces to the Hamiltonian XXZ density [Eq. \(2.45\)](#), whereas for finite λ one obtains a genuinely discrete-time integrable dynamics.

In the easy-axis regime ($J < J_z$) and at half filling, the spin Drude weight vanishes and spin transport is normal diffusive at finite temperature, while in the easy-plane regime ($J > J_z$) quasi-local charges lead to a finite Drude weight and ballistic transport. The isotropic point $J = J_z$ is dynamically special: despite integrability, spin spreading is superdiffusive with KPZ scaling, as experimentally shown on a quantum computer in [Ref. \[92\]](#). These regimes, well known in the Hamiltonian chain, persist in the integrable circuit setting, while finite step size can also generate Floquet-specific crossovers and sharp dynamical features absent in continuous time.

Because the dynamics proceeds in two substeps (odd then even bonds), charge conservation naturally yields two staggered continuity relations: one for even sites and one for odd sites, involving distinct local current operators. In particular, the “odd” current is generically a four-site operator (reflecting the two-layer update), and only reduces to the familiar two-site Hamiltonian current in the continuous-time limit. This even/odd structure is a useful book-keeping device when defining transport diagnostics in discrete time, and we return to it in [Ch. 3](#).

Recently [\[93\]](#), it has been shown that circuits composed of copies of the same $U(1)$ -

invariant two-qubit gate are integrable. We examine in more detail the structure of the conserved quantities that make such a family of circuits integrable in [Sec. 3.7.1](#), and we will use these charges as a benchmark for the prethermal decay analysis of [Sec. 3.7](#).

2.5.1 Dual-unitary circuits

Along the line $J = \pi$ the local gates of the $U(1)$ -symmetric XXZ family acquire dual unitary (DU). In this limit the two-site gate simplifies to a SWAP-type permutation (up to phases within each $\hat{S}^z$ sector). Setting $\phi = h = h' = 0$ for clarity, one finds in the computational basis

$$\hat{u}(J = \pi) = \begin{pmatrix} e^{-iJ_z/4} & 0 & 0 & 0 \\ 0 & 0 & -i e^{+iJ_z/4} & 0 \\ 0 & -i e^{+iJ_z/4} & 0 & 0 \\ 0 & 0 & 0 & e^{-iJ_z/4} \end{pmatrix}. \quad (2.49)$$

so that spin excitations propagate ballistically by successive (phase-dressed) swaps. This “SWAP point” has no continuous-time local analogue and plays a central role in the disordered model of [Sec. 2.6](#).

More generally, a two-qubit gate $\hat{u}$ is dual-unitary if it is unitary both as a time-evolution map and as a space-evolution map. Writing

$$\hat{u} = \sum_{i,j,k,l=0}^1 u_{ij}^{kl} |kl\rangle \langle ij|, \quad u_{ij}^{kl} = \langle kl | \hat{u} | ij \rangle, \quad (2.50)$$

ordinary unitarity is

$$\sum_{k,l} u_{ij}^{kl} (u_{i'j'}^{kl})^* = \delta_{i,i'} \delta_{j,j'}, \quad (2.51)$$

while the dual gate $\hat{\tilde{u}}$ is defined by reshuffling indices, $\tilde{u}_{ik}^{jl} = u_{ij}^{kl}$, and dual unitarity requires $\hat{\tilde{u}}$ to be unitary as well:

$$\sum_{j,l} u_{ij}^{kl} (u_{i'j'}^{kl})^* = \delta_{i,i'} \delta_{k,k'}. \quad (2.52)$$

Dual-unitary circuits form a broad class of exactly tractable Floquet many-body dynamics: the space-time transfer matrices are unitary, and many spatio-temporal correlation functions reduce to finite-dimensional transfer-matrix problems. In the absence of additional conservation laws, this implies particularly sharp light-cone

structures for correlations and entanglement growth. In the present $U(1)$ -symmetric XXZ setting, however, dual unitarity coexists with charge conservation, and the resulting dynamics retains nontrivial transport structure; the SWAP-type simplification Eq. (2.49) provides an analytically controlled benchmark for the more general (non-dual and disordered) regimes explored later in the thesis.

2.6 Discrete-time disordered XXZ chain




Figure 2.5: Schematics of model phases, dynamics, and a typical circuit realisation. (a) Parameter-space overview of the generic $U(1)$ -symmetric circuit: along the three solid borders the model is exactly solvable, while the dashed cut at $J_z = \pi$ traverses four distinct dynamical regimes that we use as a one-parameter slice through the phase diagram. (b) Representative initial (dotted) and infinite-time (solid) spin-magnetisation profiles along $J_z = \pi$. (c) One Floquet period is generated by a $U(1)$ -symmetric unitary $\hat{U}$ built from nearest-neighbour gates.

Beyond serving as a clean testbed for digital XXZ dynamics, the circuit of Sec. 2.5 admits a natural disordered generalisation. In continuous time, the XXZ and isotropic Heisenberg chains, originally introduced as idealised models [94], are standard workhorses for spin transport [65, 95, 96] and many-body localisation (MBL) [97–107]. Quenched disorder does not merely perturb the clean dynamics: it can stabilise regimes with strongly suppressed transport that are absent in translationally invariant chains. Discrete-time counterparts inherit the same $U(1)$ conservation law and locality [108–110], but in Floquet settings the absence of energy conservation and the presence of special points such as dual-unitary lines can generate additional, intrinsically digital, dynamical structure. In particular, in the vicinity of the DU/SWAP line discussed in Sec. 2.5.1, in Ref. [2] we identify a prethermal regime (“swappy” dynamics) with coherent propagation of excitations despite strong disorder.

We choose the disorder ensemble so as to break all symmetries except the global $U(1)$

charge. Energy conservation is absent by construction in discrete time, and translation invariance is destroyed both by random gate parameters and by randomising the update order. In this section we use this disordered circuit as the main platform for studying transport in digital matter: static, infinite-time diagnostics are summarised in [Sec. 2.6.2](#), while dynamical transport diagnostics are developed in [Ch. 3](#).

We consider N qubits evolving under repeated application of a Floquet unitary $\hat{U}$. One Floquet period is defined as a product of N nearest-neighbour gates,

$$\hat{U} = \prod_{\{P(n)\}} \hat{u}_{n,n+1}, \quad (2.53)$$

where P is a random permutation of the bond indices n (with periodic boundary conditions, so $n = N$ denotes the bond $(N, 1)$). A *single disorder realisation* consists of the choice of P together with a set of gate parameters (defined below); $\hat{U}$ is then held fixed in time and the discrete-time dynamics is $\hat{U}^t$.

The two-qubit gates are taken as

$$\hat{u}_{n,n+1} = e^{-i\hat{h}_{n,n+1}} e^{-i\hat{h}_n \hat{s}_n^z} e^{-i\hat{h}'_n \hat{s}_{n+1}^z}, \quad (2.54)$$

with an XXZ-like interaction like [Eq. \(2.45\)](#) dressed by random parameters:

$$\hat{h}_{n,n+1} = \frac{J}{2} (\hat{s}_n^+ \hat{s}_{n+1}^- e^{i\phi_n} + \hat{s}_n^- \hat{s}_{n+1}^+ e^{-i\phi_n}) + J_z \hat{s}_n^z \hat{s}_{n+1}^z, \quad (2.55)$$

where ϕ_n is a Peierls phase [\[111\]](#). Each gate preserves the two-site magnetisation, $[\hat{u}_{n,n+1}, \hat{s}_n^z + \hat{s}_{n+1}^z] = 0$, and hence the full Floquet unitary conserves the global charge

$$\hat{S}^z = \sum_{n=1}^N \hat{s}_n^z, \quad [\hat{U}, \hat{S}^z] = 0, \quad (2.56)$$

which is the prerequisite for defining spin transport.

The form [Eq. \(2.54\)](#) is not ad hoc: any two-qubit unitary commuting with $\hat{s}_n^z + \hat{s}_{n+1}^z$ is block diagonal in the $\hat{S}^z$ sectors. In the computational basis $\{|\uparrow\uparrow\rangle, |\uparrow\downarrow\rangle, |\downarrow\uparrow\rangle, |\downarrow\downarrow\rangle\}$ one has a $1 \oplus 2 \oplus 1$ decomposition, i.e. two independent phases on the $S^z = \pm 1$ sectors and a general $U(2)$ action on the single-excitation subspace. Up to an irrelevant global phase, this gives five independent parameters, which are precisely provided by $\{J, J_z, h_n, h'_n, \phi_n\}$.

Disorder is introduced by sampling the three phase parameters $\{h_n, h'_n, \phi_n\}$ independently and uniformly from $(-\pi, \pi]$ for each gate, while keeping J and J_z fixed across all gates in a given Floquet unitary. Using the $SU(4)$ Cartan (Weyl-chamber) classification of two-qubit gates up to local unitaries and SWAP [\[112\]](#), it is sufficient here to restrict

the interaction parameters in $\hat{u}_{n,n+1}$ to $J, J_z \in [0, \pi]$.

Similar U(1)-symmetric Floquet chains have appeared in a number of contexts [113–122]. Many works adopt a depth-two brickwork geometry [119–122], whereas the random update order used here is closer in spirit to Ref. [115] and treats odd and even N on the same footing. In Sec. B we provide an explicit dictionary between equivalent representations of a generic two-qubit U(1)-symmetric gate: starting from the phase-dressed XXZ form Eqs. (2.54) and (2.55), we show how the four local Z -rotation phases can be reduced to two onsite fields and a single Peierls phase (yielding a compilation-friendly decomposition for digital simulation), and we give the corresponding change of variables to the integrability-oriented R -matrix parametrisation used in [93].

2.6.1 Overview of regimes of transport dynamics

Before presenting numerical results, it is useful to highlight three exactly solvable limits, corresponding to the three solid borders of the $(J, J_z) \in [0, \pi]^2$ square in Fig. 2.5(a):

1. *Free-fermion line* ($J_z = 0$). Under Jordan-Wigner, the circuit maps to a quadratic fermion Floquet dynamics; with quenched phases $\{h_n, h'_n, \phi_n\}$ one expects Anderson localisation of single-particle modes and strongly suppressed transport.
2. *Ising line* ($J = 0$). The gates are diagonal in the $\hat{s}_n^z$ basis, so each $\hat{s}_n^z$ is conserved and spin transport is strictly absent.
3. *Dual-unitary / SWAP line* ($J = \pi$). Along this line the local gates are dual unitary (DU); in the present U(1)-symmetric family they reduce to generalised SWAPs (see Sec. 2.5.1 for the DU definition and for the clean-gate form Eq. (2.49)). For generic phases one may write

$$\hat{u}_{n,n+1}(J = \pi) = \begin{pmatrix} e^{-i\kappa_+} & 0 & 0 & 0 \\ 0 & 0 & e^{-i\xi_-} & 0 \\ 0 & e^{-i\xi_+} & 0 & 0 \\ 0 & 0 & 0 & e^{-i\kappa_-} \end{pmatrix} \quad (2.57)$$

with $\kappa_{\pm} = \mp(h + h')/2 + J_z/4$ and $\xi_{\pm} = \pi \pm (h - h')/2 - J_z/4 \mp \phi$. Thus, excitations are permuted deterministically by successive swaps and acquire only phase dressing: transport is ballistic and the dynamics is effectively noninteracting even in the presence of strong disorder.

The existence of these non-thermalising borders raises natural questions about nearby dynamics. Near the Ising line, small but nonzero J competes with strong on-site disorder and one expects strongly suppressed transport, in loose analogy with phenomenological finite-size signatures of MBL in Hamiltonian chains [99, 123]. More unusually, the DU/SWAP line has no continuous-time analogue: it is a *Floquet* integrable limit whose vicinity can host long-lived prethermal dynamics. A central outcome of Ref. [2] is that, near $J = \pi$, transport can become anomalously fast for parametrically long times (the “swappy” regime), despite the presence of strong disorder.

To probe the interpolation between the Ising and DU limits while staying far from the free-fermion line, we focus on the cut $J_z = \pi$. Along this line, four representative points are used throughout the thesis: localised ($J = 0.395$; ●), ergodic ($J = 1.374$; ►), swappy ($J = 2.551$; ■), and near-SWAP ($J = 3.138$; ★), as indicated in Fig. 2.5(b). At the system sizes accessible to exact numerics ($N \leq 22$), these regimes exhibit either complete localisation / subdiffusive spreading, diffusion, or superdiffusive spreading augmented by ballistic single-excitation propagation (see Figs. 3.3 and 3.4 and Secs. 3.5 and 3.6).

We emphasise that determining the thermodynamic stability of localisation-like behaviour in disordered Floquet systems is beyond the scope of this thesis. As in the Hamiltonian MBL literature, finite-size signatures can drift strongly with N and may be destabilised by rare-region (avalanche) mechanisms [123, 124]. Accordingly, we restrict our claims to phenomenological transport behaviour at intermediate sizes and times, with a particular focus on the nature and robustness of the swappy regime.

To characterise the model we use complementary static and dynamical diagnostics. Static spectral quantities probe infinite-time structure (Sec. 2.6.2), while dynamical observables resolve intermediate-time transport regimes that can be invisible in static chaos indicators (Ch. 3).

2.6.2 Spectral properties of a discrete-time model

Static diagnostics are extracted from the eigenphases and eigenvectors of the Floquet unitary,

$$\hat{U} |\varphi_i\rangle = e^{-i\varphi_i} |\varphi_i\rangle, \quad \varphi_i \in (-\pi, \pi]. \quad (2.58)$$

We focus on two standard indicators of ergodicity in disordered many-body systems: the mean eigenstate entanglement entropy and the mean adjacent gap ratio [123]. Both require partial diagonalisation of $\hat{U}$, which we perform using POLFED [126] (implementation details are given in Sec. 2.9.1). Definitions and benchmarks for the




Figure 2.6: Static properties of the model. (a) Mean eigenstate von Neumann entanglement entropy along $J_z = \pi$, normalised by the symmetry-resolved Page value [85], for several system sizes. Error bars show standard deviations over at least 40 disorder realisations, with outliers removed by a $3.5\times$ interquartile-range filter [125]. Residual fluctuations near the SWAP point arise from near-degeneracies in the Floquet spectrum. (b) Gap-ratio phase diagram at $N = 15$, averaged over 10–20 realisations in the fixed- $M = 0$ sector, computed with the polynomially filtered exact diagonalisation method (POLFED) eigensolver [126].

entropy (including the symmetry-resolved Page value) and the gap ratio are discussed in Sec. 2.1.3 and Sec. 2.1.1, respectively.

All spectral quantities are computed within a fixed U(1) sector, i.e. for fixed total magnetisation $\hat{S}^z |\psi\rangle = M |\psi\rangle$ (in practice we use half filling, $M = 0$), and are then averaged over disorder realisations.

We normalise the mean eigenstate entanglement entropy by the Page entropy in the same charge sector, $\langle s \rangle / s_{\text{Page}}$ (see Sec. 2.1.3), and discuss efficient evaluation in Sec. 2.9.2. As shown in Fig. 2.6(a), $\langle s \rangle / s_{\text{Page}}$ is small in the strongly localised regime ($J \ll \pi$) and approaches unity in the ergodic regime, consistent with thermal eigenstates at accessible sizes. Apparent deviations from the Page benchmark close to the SWAP point $J \rightarrow \pi$ shrink with increasing N . This is consistent with the fact that the DU/SWAP limit carries additional (exact or approximate) conservation laws beyond $\hat{S}^z$ (see Sec. 2.5.1), which can fragment the finite-size spectrum into overlapping subspectra and distort raw chaos indicators if those sectors are not first resolved. We return to this mechanism, from a transfer-matrix viewpoint, in Sec. D.

The second diagnostic is the mean adjacent gap ratio $\langle r \rangle$, computed from ordered eigenphases within a fixed- M sector as defined in Sec. 2.1.1. Because the disorder ensemble includes generic complex phases (notably ϕ_n), it generically breaks time-reversal symmetry; accordingly, the ergodic benchmark is the circular unitary ensemble, $\langle r \rangle \simeq r_{\text{CUE}} \approx 0.6027$, whereas integrable/localised spectra yield

$\langle r \rangle \simeq r_{\text{Poi}} \approx 0.3863$ [72]. As shown in Fig. 2.6(b), $\langle r \rangle$ interpolates from Poisson-like values at small J to CUE-like values in the ergodic regime, with clear integrable signatures on the free-fermion line $J_z = 0$ and on the dual-unitary line $J = \pi$.

Taken together, $\langle s \rangle$ and $\langle r \rangle$ probe complementary aspects of the infinite-time structure. The entropy tests whether typical Floquet eigenstates have thermal reduced states (as expected if a Floquet-ETH holds), while the gap ratio diagnoses spectral correlations (level repulsion versus clustering). However, these are *asymptotic* diagnostics: they do not by themselves resolve long-lived prethermal regimes. In particular, the swappy regime identified dynamically near $J = \pi$ in Ref. [2] is largely invisible to these static measures at accessible sizes, motivating the transport diagnostics developed in Ch. 3.

Finally, we note that the apparent crossing points extracted from $\langle r \rangle$ and from $\langle s \rangle / s_{\text{Page}}$ drift systematically with N over the accessible range. Such finite-size drifts are ubiquitous in ED studies of disordered interacting chains and were historically used as practical probes of MBL-like crossovers [71]. More recent work cautions that extrapolation is subtle: rare thermal regions and avalanche-type instabilities can destroy apparent localisation at larger scales [127, 128], and slow crossovers and strong sample-to-sample fluctuations complicate direct interpretation [129, 130]. Accordingly, while our data are consistent with a crossover from ergodic to strongly suppressed-transport behaviour at accessible N , they do not constitute evidence for a stable MBL phase in the thermodynamic limit.

2.7 Current operators and the single-impurity circuit

In this section we define the current operators for continuous and discrete time systems and introduce a benchmark in which a strictly *local* perturbation strongly affects spectral chaos diagnostics while leaving infinite-time transport essentially unchanged. The resulting “impurity test” can be applied both in continuous time and in discrete-time circuits. Its purpose is methodological: it separates genuinely hydrodynamic statements (scaling of the zero-frequency weight of current correlations) from finite-size indicators of chaos (level statistics and eigenstate structure). Concretely, it connects the projector formulation of infinite-time (Cesàro) averages from Sec. 2.1 with the continuity-equation construction of currents and Drude diagnostics from Sec. 2.3. It also serves as a recurring cautionary example when comparing static chaos indicators (e.g. gap ratios in Sec. 2.6.2) to transport diagnostics, a theme we revisit for the clean circuit in Sec. 2.5 and the disordered model in Sec. 2.6.

2.7.1 Continuous time

Following Refs. [131, 132], the spin- $\frac{1}{2}$ XXZ chain with a single magnetic defect near the centre has become a canonical benchmark for disentangling three notions that are often conflated: integrability, ETH, and transport. The point is that a strictly local perturbation can induce Wigner-Dyson level statistics and enforce ETH-like structure of matrix elements for local observables, while leaving the bulk U(1) hydrodynamics, and in particular the ballistic channel in the easy-plane regime, essentially unchanged in the thermodynamic limit.

The unperturbed Hamiltonian is

$$\hat{H}_{\text{XXZ}} = \sum_{n=1}^N \left(J \hat{s}_n^x \hat{s}_{n+1}^x + J \hat{s}_n^y \hat{s}_{n+1}^y + J_z \hat{s}_n^z \hat{s}_{n+1}^z \right), \quad (2.59)$$

and the impurity model adds a local longitudinal field (for even N , at $n_0 = N/2$),

$$\hat{H}_{\text{SI}} = \hat{H}_{\text{XXZ}} + h \hat{s}_{n_0}^z. \quad (2.60)$$

Both conserve $\hat{S}^z = \sum_n \hat{s}_n^z$ and thus support a well-defined magnetisation current.

The current is fixed by the local continuity equation for the conserved density $\hat{s}_n^z$. Away from the impurity site the defect term commutes with $\hat{s}_n^z$ and does not modify the continuity equation Eq. (2.29), so the standard XXZ bond current density applies,

$$\hat{j}_n^z = J(\hat{s}_n^x \hat{s}_{n+1}^y - \hat{s}_n^y \hat{s}_{n+1}^x), \quad \hat{J}^z = \sum_{n=1}^N \hat{j}_n^z. \quad (2.61)$$

Transport is then diagnosed by the current autocorrelator and its spectral density, as introduced in Sec. 2.3.

To make explicit what is (and is not) controlled by ETH, it is useful to separate diagonal and off-diagonal contributions in the energy basis, using the projection identities of Sec. 2.1. For an initial state $|\psi(0)\rangle = \sum_i c_i |E_i\rangle$ the density matrix decomposes as

$$\hat{\rho}(t) = \sum_i |c_i|^2 |E_i\rangle\langle E_i| + \sum_{i \neq j} c_i c_j^* e^{-it(E_i - E_j)} |E_i\rangle\langle E_j| = \hat{\rho}_{\text{diag}} + \hat{\rho}_{\text{off}}(t), \quad (2.62)$$

so that the infinite-time (Cesàro) average of any observable reads

$$\overline{A} = \text{Tr}[\hat{\rho}_{\text{diag}} \hat{A}], \quad (2.63)$$

while $\hat{\rho}_{\text{off}}(t)$ controls temporal fluctuations and, crucially for transport, the frequency-resolved spectral weight of correlators.

In an ergodic system satisfying ETH (Sec. 2.1.2), the diagonal matrix elements A_{ii} become a smooth function of energy, so $\hat{\rho}_{\text{diag}}$ reproduces the appropriate equilibrium

value for a broad class of narrow-band initial states; conversely, in integrable systems the failure of diagonal ETH implies that $\hat{\rho}_{\text{diag}}$ generally retains memory of additional conserved structures (captured by GGEs).

Equivalently, at the level of matrix elements one may write

$$\langle \hat{A}(t) \rangle = \sum_{i,j} c_i^* c_j e^{i(E_i - E_j)t} A_{ij}. \quad (2.64)$$

For a non-degenerate spectrum, the infinite-time (Cesàro) average removes the oscillatory terms and yields the diagonal ensemble Eq. (2.63), while the explicitly time-dependent piece is controlled by the off-diagonal matrix elements $A_{i \neq j}$ and governs temporal fluctuations and linear-response spectra. The single-impurity benchmark is striking because the impurity enforces ETH-like *diagonal* and *off-diagonal* statistics, yet the low-frequency structure of the current matrix elements, and therefore the ballistic contribution in Eq. (2.36), can remain essentially the one inherited from the integrable bulk.

The off-diagonal sector matters even more directly for linear-response transport. For the spin conductivity one may write the Kubo decomposition

$$\text{Re } \sigma_N(\omega) = 2\pi D_N \delta(\omega) + \frac{\pi}{N} \frac{1 - e^{-\beta\omega}}{\omega} \sum_{i \neq j} p_i |J_{ij}^z|^2 \delta(E_j - E_i - \omega), \quad p_i = \frac{e^{-\beta E_i}}{Z}, \quad (2.65)$$

which makes explicit that the regular low-frequency response is governed by off-diagonal current matrix elements. In the present benchmark this is the key organising principle: the impurity primarily reshuffles eigenvectors and the statistics of matrix elements (diagonal and off-diagonal ETH), whereas the $\omega \rightarrow 0$ envelope of $|J_{ij}^z|^2$ that controls transport can remain essentially unchanged.

This is precisely what is observed in the easy-plane regime $0 < J_z < 1$. Ref. [131] shows that the impurity makes the spectrum quantum chaotic in the usual random-matrix sense, while high-temperature spin transport remains ballistic. From the ensemble viewpoint, the single impurity restores diagonal ETH for local observables away from the defect (the spread of A_{ii} within a microcanonical window shrinks with N), and it renders the off-diagonal matrix elements consistent with the ETH form of Sec. 2.1.2; remarkably, for current operators the low-frequency structure that governs transport is inherited from the integrable XXZ chain [132]. As mentioned in Part I, in finite open chains ballistic transport appears as a pronounced peak (or peaks) in $\text{Re } \sigma_N(\omega)$ at frequencies $\omega \sim 1/N$ whose integrated weight stays finite; increasing N shifts this weight towards $\omega \rightarrow 0$, yielding a ballistic channel in the thermodynamic limit even though the system is quantum chaotic in its level statistics. This benchmark motivates a stance that recurs below: for finite-size data, the robust quantity is the *integrated*

low-frequency weight of current correlations, while the distinction between a strict $\delta(\omega)$ contribution and a narrow peak is operational and can depend on boundary conditions and finite- N symmetry resolution (see [Secs. 2.1](#) and [2.3](#)).

2.7.2 Discrete time

We now formulate the discrete-time analogue. The role of $e^{-i\hat{H}t}$ is played by the one-period unitary $\hat{U}$, and transport quantities are defined stroboscopically. The conserved density is again $\hat{s}_n^z$, and “conductivity” refers to the linear response of the conserved current to a weak static bias, encoded in the low-frequency behaviour of the stroboscopic current correlator and its 2π -periodic spectral function.

The impurity mechanism is implemented by choosing a bulk unitary $\hat{U}_{\text{bulk}}$ with integrable structure (for instance the integrable Floquet XXZ circuit of [Sec. 2.5](#)), and inserting a strictly local modification that preserves the $U(1)$ charge. We write

$$\hat{U}_{\text{SI}} = \hat{U}_{\text{bulk}} \hat{u}_{\text{imp}}, \quad (2.66)$$

where $\hat{u}_{\text{imp}}$ acts on $\mathcal{O}(1)$ sites and commutes with $\hat{S}^z$. The discrete-time benchmark question is the direct analogue of the Hamiltonian case: does a local impurity eliminate the ballistic component, or does it merely reorganise spectral correlations while leaving the integrated low-frequency weight stable as $N \rightarrow \infty$?

To define the current operators, we start from the local continuity equation for the conserved density, same as [Eq. \(2.30\)](#). In discrete time, the one-period increment of a local density $\hat{s}_n^z$ is expressed in terms of bond currents as

$$\hat{s}_n^z(t+1) - \hat{s}_n^z(t) = \hat{j}_{n-1}^z(t) - \hat{j}_n^z(t), \quad \hat{j}_n^z(t) = \hat{U}^{-t} \hat{j}_n^z \hat{U}^t. \quad (2.67)$$

To construct the bond currents $\hat{j}_n^z$, we first define the current associated to a single two-site gate. Let $\hat{u}_{n,n+1}$ be a two-site gate conserving the two-site magnetisation, $[\hat{u}_{n,n+1}, \hat{s}_n^z + \hat{s}_{n+1}^z] = 0$. The gate redistributes charge between the two sites while preserving the sum, so the associated right-moving bond flux is naturally defined by the increment of the right density,

$$\hat{j}_n^z[\hat{u}_{n,n+1}] = \hat{u}_{n,n+1}^\dagger \hat{s}_{n+1}^z \hat{u}_{n,n+1} - \hat{s}_{n+1}^z = -\left(\hat{u}_{n,n+1}^\dagger \hat{s}_n^z \hat{u}_{n,n+1} - \hat{s}_n^z\right). \quad (2.68)$$

For a one-period Floquet update specified by a permutation P of the nearest-neighbour bonds, as in [Eq. \(2.53\)](#), it is convenient to label the ordered gates as $\hat{u}_k = \hat{u}_{P(k), P(k)+1}$ with $k = 1, \dots, N$ (where $P(k) + 1$ is understood modulo N for periodic boundaries).

Therefore the unitary can be expressed as

$$\hat{U} = \prod_{\{P(n)\}} \hat{u}_{n,n+1} = \prod_{k=1}^N \hat{u}_k = \hat{u}_N \cdots \hat{u}_2 \hat{u}_1, \quad (2.69)$$

and we can define the partial products $\hat{U}_{<k} = \hat{u}_{k-1} \cdots \hat{u}_1$ with $\hat{U}_{<1} = \hat{I}$. A standard telescopic expansion then gives, for any local $\hat{s}_n^z$,

$$\hat{U}^\dagger \hat{s}_n^z \hat{U} - \hat{s}_n^z = \sum_{k=1}^N \hat{U}_{<k}^\dagger \left(\hat{u}_k^\dagger \hat{s}_n^z \hat{u}_k - \hat{s}_n^z \right) \hat{U}_{<k}. \quad (2.70)$$

By strict locality, the summand vanishes unless $\hat{u}_k$ acts on a bond containing site n . For the usual architecture in which each bond is applied exactly once per period, there are only two such gates: the “right” bond $(n, n+1)$ and the “left” bond $(n-1, n)$. Denoting their positions in the ordered product by $k_R = P^{-1}(n)$ and $k_L = P^{-1}(n-1)$, respectively, Eq. (2.70) reduces to two contributions,

$$\hat{U}^\dagger \hat{s}_n^z \hat{U} - \hat{s}_n^z = \hat{U}_{<k_L}^\dagger \left(\hat{u}_{n-1,n}^\dagger \hat{s}_n^z \hat{u}_{n-1,n} - \hat{s}_n^z \right) \hat{U}_{<k_L} + \hat{U}_{<k_R}^\dagger \left(\hat{u}_{n,n+1}^\dagger \hat{s}_n^z \hat{u}_{n,n+1} - \hat{s}_n^z \right) \hat{U}_{<k_R}. \quad (2.71)$$

This motivates defining the dressed bond currents at a chosen stroboscopic reference time by dressing the gate-local current operators $\hat{j}^z[\hat{u}]$ with the appropriate prefix unitary. Concretely,

$$\hat{j}_n^z = \hat{U}_{<k_R}^\dagger \hat{j}_n^z[\hat{u}_{n,n+1}] \hat{U}_{<k_R}, \quad \hat{j}_{n-1}^z = \hat{U}_{<k_L}^\dagger \hat{j}_{n-1}^z[\hat{u}_{n-1,n}] \hat{U}_{<k_L}. \quad (2.72)$$

With these definitions, the one-period increment obeys the discrete continuity equation,

$$\hat{s}_n^z(t+1) - \hat{s}_n^z(t) = \hat{j}_{n-1}^z(t) - \hat{j}_n^z(t), \quad \hat{j}_n^z(t) = \hat{U}^{-t} \hat{j}_n^z \hat{U}^t. \quad (2.73)$$

The extensive current entering linear response is

$$\hat{J}^z = \sum_{n=1}^N \hat{j}_n^z, \quad \hat{J}^z(t) = \hat{U}^{-t} \hat{J}^z \hat{U}^t. \quad (2.74)$$

The permutation P affects the microscopic support of $\hat{j}_n^z$ through the dressing unitaries $\hat{U}_{<k}$. Nevertheless, stroboscopic transport coefficients built from $\hat{J}^z(t)$ are invariant under these choices.

All transport diagnostics introduced earlier carry over with the replacement $t \in \mathbb{R} \mapsto t \in \mathbb{Z}$. The stroboscopic current correlator in an ensemble $\langle \cdot \rangle$ (here: infinite temperature within a fixed $\hat{S}^z$ sector) is

$$C_J(t) = \frac{1}{N} \langle \hat{J}^z(t) \hat{J}^z(0) \rangle, \quad t \in \mathbb{Z}, \quad (2.75)$$

and its spectral function is the 2π -periodic Fourier series

$$S_J(\omega) = \sum_{t \in \mathbb{Z}} e^{i\omega t} C_J(t), \quad \omega \in [-\pi, \pi]. \quad (2.76)$$

The stroboscopic Drude weight is extracted from the infinite-time Cesàro average,

$$D \propto \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} C_J(t). \quad (2.77)$$

For finite-size numerics it is convenient to evaluate the Cesàro limit directly from the spectrum. Let $\{\hat{P}_\alpha\}$ denote the spectral projectors of $\hat{U}$ within the chosen $\hat{S}^z$ sector (projectors onto degenerate eigenphase subspaces). Then the time average is a projection onto the commutant of $\hat{U}$ and one finds

$$\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} C_J(t) = \frac{1}{N} \frac{1}{D_M} \sum_{\alpha} \text{Tr}[\hat{P}_\alpha \hat{J}^z \hat{P}_\alpha \hat{J}^z], \quad (2.78)$$

where $D_M = \dim \mathcal{H}^{(M)}$. In the absence of degeneracies this reduces to a sum over Floquet eigenstates,

$$\overline{C_J} = \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} C_J(t) = \frac{1}{N} \frac{1}{D_M} \sum_{i=1}^{D_M} |\langle \varphi_i | \hat{J}^z | \varphi_i \rangle|^2, \quad (2.79)$$

Eqs. (2.78) and (2.79) are our default estimators of the finite-size Drude weight throughout the thesis.

2.7.3 Numerical protocol and results

We implement the single-impurity benchmark in discrete time with open boundary condition (OBC) by perturbing the integrable U(1)-symmetric Floquet XXZ circuit $\hat{U}_{\text{bulk}}$ of Sec. 2.5 with a defect supported on a single bond or a single site. We perturb the clean model with an on-site kick $\hat{u}_{\text{imp}} = \exp(-i h_{\text{imp}} \hat{s}_{n_0}^z)$ applied once per Floquet period. This preserves the conserved magnetisation $\hat{S}^z$, but destroys the commuting-transfer-matrix structure that renders the bulk circuit integrable, as mentioned in Sec. 2.5.

All computations are performed at half filling, $M = 0$ sector of dimension $D_M = \binom{N}{N/2+M}$. For each system size N we diagonalise the perturbed Floquet unitary $\hat{U}_{\text{SI}}$ within this sector, obtaining eigenpairs $\hat{U}_{\text{SI}}|\varphi_i\rangle = e^{-i\varphi_i}|\varphi_i\rangle$. Spectral chaos is quantified by the mean adjacent gap ratio $\langle r \rangle$ computed from the ordered eigenphases $\{\varphi_i\}$.

To probe transport we use the sector-projected spin current operator $\hat{J}^z$ and evaluate




Figure 2.7: Single-impurity benchmark in discrete time with OBC at half filling along the fixed-coupling ray $J_z = 0.55J$, with $h_{\text{imp}} = 0.5J$ and $h_{\text{edge}} = 0.1J$. (a) Left: mean adjacent gap ratio $\langle r \rangle$ from Floquet eigenphases, showing a fast crossover from Poisson-like to COE-like statistics once a local defect is present. (b) Right: finite-size Drude estimator $\overline{C_J}$ computed from Eq. (2.79). The key point is that level statistics are strongly modified by the impurity, while the intensive zero-frequency weight controlling ballistic transport is comparatively stable.




Figure 2.8: Regular (off-diagonal) current spectral density $S_J^{\text{reg}}(\omega)$ for OBC, estimated from Floquet current matrix elements using Eq. (2.80) (a frequency-binned sum of $|\langle \varphi_i | \hat{J}^z | \varphi_j \rangle|^2$ over pairs with $i \neq j$ in the infinite-temperature, fixed- M ensemble), with eigenphase differences wrapped to $[-\pi, \pi)$ and bin width $\delta\omega = 0.025J$. Spectra are displayed as functions of $\omega N/J$ to remove the trivial kick-dependent frequency scale and to emphasise the finite-size sharpening of the low-frequency structure. The impurity reorganises detailed finite-size features across ω but does not suppress the integrated low-frequency weight, consistent with the Drude estimates in Fig. 2.7.

two complementary diagnostics. First, the ballistic channel is extracted from the diagonal ensemble associated with stroboscopic time averaging. In the absence of eigenphase degeneracies, the infinite-time average of the current autocorrelator reduces to

$$\overline{C_J} = \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} C_J(t) = \frac{1}{N} \frac{1}{D_M} \sum_{i=1}^{D_M} |\langle \varphi_i | \hat{J}^z | \varphi_i \rangle|^2, \quad (2.79)$$

which provides our finite-size Drude estimator.² With $\hat{J}^z = \sum_n \hat{j}_n^z$ extensive, the prefactor $1/N$ makes Eq. (2.79) intensive.

Second, we compute a frequency-resolved *regular* spectrum from the off-diagonal matrix elements of $\hat{J}^z$ in the Floquet eigenbasis. Eigenphase differences are wrapped into the symmetric interval $[-\pi, \pi)$ to avoid folding negative frequencies onto values near 2π , and binned with width $\delta\omega$. The resulting binned spectrum is normalised as

$$S_J^{\text{reg}}(\omega_k) \approx \frac{1}{N D_M \delta\omega} \sum_{i \neq j: \Delta\varphi_{ij} \in \text{bin } k} |\langle \varphi_i | \hat{J}^z | \varphi_j \rangle|^2, \quad (2.80)$$

This is a discretised estimator of the Lehmann spectral density $S_J(\omega)$ in the infinite-temperature, fixed- M ensemble, restricted to off-diagonal pairs (hence “regular”). If one defines the infinite-temperature regular conductivity by $\sigma_\infty(\omega) := \lim_{\beta \rightarrow 0} \text{Re } \sigma_{\text{reg}}(\omega)/\beta$, then Eq. (2.34) gives $\sigma_\infty(\omega) = \pi S_J^{\text{reg}}(\omega)$ for $\omega \neq 0$. So that $\sum_k S_J^{\text{reg}}(\omega_k) \delta\omega = (N D_M)^{-1} \sum_{i \neq j} |\langle \varphi_i | \hat{J}^z | \varphi_j \rangle|^2$ is a finite-size Hilbert-Schmidt sum rule. By construction, Eq. (2.80) omits $i=j$ and therefore captures only the regular part; the ballistic weight is assessed separately through Eq. (2.79).

The single impurity is implemented as an additional local Z -kick applied once per Floquet period at the centre of the chain. Let $n_0 = \lfloor \frac{N}{2} \rfloor$. We place the impurity on the bond $(n_0, n_0 + 1)$ by setting $h_{n_0} = h_{\text{imp}}$ (and $h_{n_0+1} = 0$), so that $\hat{u}_{n_0, n_0+1} = \hat{R}_{n_0}^z(h_{\text{imp}}) \hat{u}_{\text{XXZ}}(J, J_z)$, with $\hat{u}_{\text{XXZ}}(J, J_z)$ the integrable Floquet XXZ gate of Sec. 2.5. To lift residual symmetries and accidental degeneracies, and to stabilise the non-degenerate estimator Eq. (2.79), we additionally apply an edge rotation $\hat{R}^z(h_{\text{edge}})$ on the first bond. All other gates are unmodified, $\hat{u}_{n, n+1} = \hat{u}_{\text{XXZ}}(J, J_z)$, for $n \neq n_0$. We take the impurity strength to be $h_{\text{imp}} = 0.5J$.

When scanning the kick strength J we follow a fixed ray in parameter space by also keeping the other coupling ratios constant, as the longitudinal coupling $J_z = 0.55 J$ and the single edge kick field $h_{\text{edge}} = 0.1 J$. For the same reason, the histogram resolution is chosen proportional to J , $\delta\omega = 0.025 J$, which keeps the binning resolution fixed in the scaled variable ω/J and avoids purely numerical J -dependent broadening. Accordingly, we plot the regular spectrum against $\omega N/J$, which removes the trivial kick-dependent

²If exact degeneracies are present, the basis-independent generalisation replaces the sum of diagonal terms by the Hilbert-Schmidt weight of $\hat{J}^z$ within each degenerate eigenphase subspace. Away from fine-tuned points, the diagonal expression is numerically stable.

quasienergy scale and highlights finite-size sharpening near $\omega \simeq 0$. This rescaling is a presentation device: the Drude weight Eq. (2.79) is extracted independently and is unaffected by the choice of horizontal axis.

The results are summarised in Figs. 2.7 and 2.8. A single defect drives a rapid crossover in spectral statistics from Poisson-like values (integrable bulk) to COE-like values (quantum-chaotic spectrum), visible already at modest N . The intensive ballistic weight extracted from the latter remains stable within finite-size trends: the impurity substantially reorganises eigenphases and off-diagonal matrix elements, yet does not eliminate the integrated zero-frequency weight that signals ballistic transport. The regular spectrum Eq. (2.80) exhibits strong finite-size structure and impurity-induced reshuffling across frequencies; nevertheless, the low-frequency weight does not collapse, consistent with the persistence of a nonzero Drude contribution.

Note that both Eqs. (2.79) and (2.80) can potentially be estimated with the method outlined in Sec. 2.4.2 and discussed further in Sec. 2.9.7.

2.8 Conclusion

This chapter developed a spectral language for *infinite-time* dynamics in Hamiltonian and Floquet many-body systems. By rewriting long-time limits as spectral projections (Sec. 2.1) and organising operator content through the DOS, LDOS, and related spectral functions (Sec. 2.2), we obtained a compact set of objects that connect directly to linear response and to transport coefficients. Specialising to U(1) spin transport, we used current correlators and their low-frequency structure to define Drude-weight diagnostics in continuous and stroboscopic time (Sec. 2.3), and benchmarked these tools in both Hamiltonian and circuit settings (Secs. 2.4 to 2.7).

Spectral chaos and transport are separate concepts. Level statistics and eigenstate structure diagnose ergodicity, but they do not by themselves fix the $\omega \rightarrow 0$ envelope of current matrix elements that controls the integrated low-frequency weight of response functions. The single-impurity benchmark of Sec. 2.7 makes this separation clear: a strictly local perturbation can switch the eigenphase statistics from Poissonian to random-matrix-like while leaving an extensive ballistic channel essentially inherited from an integrable bulk. The operator-resolved spectral densities therefore play a dual role in this thesis: they quantify “how chaotic” a system is, and they specify *which operators* couple to the slow modes that govern transport.

At the same time, the quantities emphasised here are asymptotic: they describe infinite-time limits or frequency-resolved weight. In Floquet circuits in particular, where experiments and numerics access only a finite number of cycles, dynamics

can be dominated by long-lived intermediate-time regimes and crossovers even when infinite-time diagnostics look featureless. In [Ch. 3](#) we turn our attention to finite-time transport probes, built from correlation functions of (quasi)-conserved quantities, that expose transient and prethermal windows (including the near-SWAP “swappy” regime) which are largely invisible to infinite-time limits alone.

2.9 Methods

We now collect the numerical and algorithmic tools used in this chapter to access *infinite-time* information. As established in [Sec. 2.1](#), taking $T \rightarrow \infty$ suppresses oscillatory contributions and reduces long-time statements to spectral data: eigenphases and eigenvectors for Floquet unitaries $\hat{U}$, or eigenvalues and eigenvectors for Hamiltonians $\hat{H}$. As the primary focus of this chapter is the analysis of Floquet systems, our methodology is tailored to address their study.

First, to diagnose ergodicity and the structure of eigenstates, namely the results presented in [Sec. 2.6](#) (Ref. [2]), we require selected eigenpairs, typically in a narrow eigenphase (or energy) window. For Floquet systems we obtain these efficiently via polynomially filtered exact diagonalisation method (POLFED), a unitary analogue of shift-invert that accelerates Arnoldi convergence in a target phase window; see [Sec. 2.9.1](#). We then evaluate symmetry-resolved eigenstate entanglement using block-wise Schmidt decompositions within fixed magnetisation sectors; see [Sec. 2.9.2](#).

Second, we move to the tools needed to access spectral densities and coarse-grained thermodynamic information beyond exact diagonalisation. We start with discussing state typicality and stochastic trace estimation in [Sec. 2.9.3](#), which we will also employ in [Ch. 3](#). For Hamiltonians this is the kernel polynomial method (KPM), which reconstructs the DOS (and operator-resolved spectral functions) from Chebyshev moments; see [Sec. 2.9.4](#).

We present a DQC1-inspired quantum procedure for moment estimation in [Sec. 2.9.5](#), which is the core of the quantum algorithmic contribution of this thesis.

We then present the results from the implementation of this algorithm on Quantinuum H1-1 hardware [Sec. 2.9.6](#).

Finally, since discrete-time dynamics is generated by $\hat{U}$ rather than $\hat{H}$, we close by formulating the corresponding moment expansions for eigenphase densities and by clarifying when Fourier reconstruction on the circle is optimal versus when a Hermitian embedding enables a Chebyshev-KPM treatment of observable-resolved spectra; see [Sec. 2.9.7](#).

2.9.1 POLFED

The eigenstate entanglement entropies reported in this chapter are computed by averaging over

$$N_{\text{ev}} = \min(d/10, 750) \quad (2.81)$$

eigenvectors in a fixed symmetry sector of dimension d . To obtain these eigenpairs efficiently for a Floquet unitary $\hat{U}$, we use the polynomially filtered exact diagonalisation method (POLFED) scheme [133], which is a polynomial spectral filter designed for unitary spectra and plays a role analogous to shift-invert for Hermitian problems.

Given a target eigenphase ϕ_{tgt} , POLFED constructs the non-unitary operator

$$\hat{K}_{M_{\text{pf}}}(\phi_{\text{tgt}}) = \sum_{k=0}^{M_{\text{pf}}} e^{-ik\phi_{\text{tgt}}} \hat{U}^k, \quad (2.82)$$

which shares eigenvectors with $\hat{U}$. Indeed, if $\hat{U}|\varphi_i\rangle = e^{-i\varphi_i}|\varphi_i\rangle$ then

$$\hat{K}_{M_{\text{pf}}}(\phi_{\text{tgt}})|\varphi_i\rangle = \kappa_{M_{\text{pf}}}(\varphi_i - \phi_{\text{tgt}})|\varphi_i\rangle, \quad \kappa_{M_{\text{pf}}}(\delta) = \sum_{k=0}^{M_{\text{pf}}} e^{-ik\delta}. \quad (2.83)$$

The filter magnitude $|\kappa_{M_{\text{pf}}}(\delta)|$ is sharply peaked at $\delta = 0$, with a main lobe of width $\Delta\varphi \sim 2\pi/(M_{\text{pf}} + 1)$. As a consequence, applying an Arnoldi iteration to $\hat{K}_{M_{\text{pf}}}$ preferentially converges to eigenvectors of $\hat{U}$ whose eigenphases lie in a narrow window around ϕ_{tgt} . Throughout this thesis we set $\phi_{\text{tgt}} = 0$. Following Refs. [126, 133], we choose the polynomial degree as

$$M_{\text{pf}} = 0.4 \frac{d}{N_{\text{ev}}}, \quad (2.84)$$

which empirically balances window selectivity against numerical conditioning and the cost of repeated applications of $\hat{U}$.

2.9.2 Efficient computation of the entropy in a magnetisation subsector

For a bipartition $A|B$, the von Neumann entanglement entropy of a pure state $|\psi\rangle$, defined in Eqs. (2.13) and (2.14), can be obtained from the Schmidt coefficients $\{s_i\}$:

$$s(|\psi\rangle) = - \sum_i s_i^2 \ln s_i^2. \quad (2.85)$$

Numerically, one reshapes the wavefunction amplitudes into a $D_A \times D_B$ matrix and computes its singular values using singular value decomposition (SVD); this step

dominates the cost.

When the dynamics conserves total magnetisation, the SVD can be performed blockwise. Let $\hat{M} = \sum_n \hat{s}_n^z$ and fix a global sector $\mathcal{H}_M$ (in this chapter we typically work at $M = 0$). The subsystem Hilbert spaces decompose as

$$\mathcal{H}_A = \bigoplus_{M_A} \mathcal{H}_{A,M_A}, \quad \mathcal{H}_B = \bigoplus_{M_B} \mathcal{H}_{B,M_B}, \quad (2.86)$$

so that the fixed-charge sector factorises as

$$\mathcal{H}_M = \bigoplus_{M_A} \mathcal{H}_{A,M_A} \otimes \mathcal{H}_{B,M-M_A}. \quad (2.87)$$

Accordingly, the reshaped coefficient matrix of $|\psi\rangle \in \mathcal{H}_M$ is block-diagonal with blocks labelled by M_A . For each M_A , we extract the corresponding component of $|\psi\rangle$, reshape it into a $d_{A,M_A} \times d_{B,M-M_A}$ matrix, compute its singular values $\{s_{M_A,i}\}$, and assemble the entropy as

$$s(|\psi\rangle) = - \sum_{M_A} \sum_i s_{M_A,i}^2 \ln s_{M_A,i}^2. \quad (2.88)$$

This replaces a single SVD at dimension $D_A \times D_B$ by a set of smaller SVDs whose largest block is parametrically smaller at fixed total charge, yielding a substantial reduction in runtime and memory footprint.

2.9.3 Stochastic trace evaluation and recursion

For overwhelming majority of pure states drawn at random from a high-dimensional subspace, expectation values concentrate sharply around ensemble averages [134, 135]. In its simplest form, typicality turns traces into averages over a small number of random probes, an idea that underlies both the classical KPM [136, 137] and a broad literature on randomised trace estimators [138, 139].

For a space $\mathcal{H}$ of dimension D with orthonormal basis $\{|i\rangle\}_{i=1}^D$, one can draw a set of probe states

$$|r\rangle = \sum_{i=1}^D c_{ri} |i\rangle, \quad (2.89)$$

with independent complex coefficients obeying

$$\mathbb{E}[c_{ri}] = 0, \quad \mathbb{E}[c_{ri}c_{rj}] = 0, \quad \mathbb{E}[c_{ri}^*c_{rj}] = \frac{\delta_{ij}}{D}, \quad (2.90)$$

so that $\mathbb{E}[|r\rangle\langle r|] = \hat{I}/D$ and therefore, for any operator $\hat{X}$ acting within the space,

$$\frac{1}{D} \text{Tr}[\hat{X}] = \mathbb{E}[\langle r| \hat{X} |r\rangle]. \quad (2.91)$$

Replacing the ensemble average by R independent probes yields the unbiased estimator

$$\frac{1}{D} \text{Tr}[\hat{X}] \approx \frac{1}{R} \sum_{r=1}^R \langle r | \hat{X} | r \rangle. \quad (2.92)$$

Typicality implies that the statistical fluctuations of Eq. (2.92) are suppressed by the *effective dimension* of the relevant ensemble (at infinite temperature, D itself), and decay at least as $R^{-1/2}$ for independent probes. For coefficient distributions with controlled fourth moments one can derive explicit variance expressions [136, 138, 139]; in many-body settings one often finds that $R = \mathcal{O}(1) - \mathcal{O}(10)$ already gives stable estimates for smooth spectral quantities because D grows exponentially with N .

Typicality is not restricted to static traces. For real-time dynamics one may approximate infinite-temperature correlation functions by evolving a *single* random state in time, a method known as dynamical quantum typicality (DQT) [60, 140, 141]. For operators $\hat{A}, \hat{B}$ (again acting within a fixed symmetry sector),

$$C_A(t) = \frac{1}{D} \text{Tr}[\hat{A}(t)\hat{A}], \quad \hat{A}(t) = e^{i\hat{H}t} \hat{A} e^{-i\hat{H}t}, \quad (2.93)$$

is typically well approximated by

$$C_A(t) \approx \langle r | \hat{A}(t) \hat{A} | r \rangle, \quad (2.94)$$

with an error controlled by D (and reduced further by averaging over R probes). At finite temperature one can combine typicality with a thermal filter to represent the Gibbs state by a *thermal typical* pure state [142]:

$$C_A^{(\beta)}(t) = \frac{\text{Tr}[e^{-\beta\hat{H}} \hat{A}(t) \hat{A}]}{Z(\beta)} \approx \frac{\langle r | e^{-\beta\hat{H}/2} \hat{A}(t) \hat{A} e^{-\beta\hat{H}/2} | r \rangle}{\langle r | e^{-\beta\hat{H}} | r \rangle}. \quad (2.95)$$

In practice, the action of $e^{-\beta\hat{H}/2}$ and the unitary time evolution can be implemented by Krylov methods, Trotterisation, or polynomial expansions; the same Chebyshev machinery used for KPM provides a natural route to approximate such filters and propagators [60, 136, 137, 143].

2.9.4 The Kernel Polynomial Method

This section introduces the KPM as an efficient method to spectral densities in many-body systems, avoiding full diagonalisation. The basic strategy is to reconstruct a target spectral function from a finite set of polynomial moments of $\hat{H}$, estimated efficiently by typicality and a stable recursion. We focus on the Chebyshev-Jackson scheme, which is well matched to sparse operators and admits controlled broadening

at fixed computational cost.

Moment expansions as spectral filters

Let $\hat{H}$ be Hermitian and rescaled so that $\text{spec}(\hat{H}) \subset (-1, 1)$. Spectral densities such as the DOS can be written as $g(E) = \frac{1}{D} \text{Tr}[\delta(E - \hat{H})]$. Replacing the sharp projector $\delta(E - \hat{H})$ by a degree- M polynomial produces a controlled spectral filter: one resolves structure on a scale set by M while avoiding exponentially expensive diagonalisation. Chebyshev polynomials are the natural basis on a bounded interval, and the KPM uses them together with kernel damping to suppress truncation-induced Gibbs oscillations [136].

Chebyshev polynomials of the first kind can be viewed as the Fourier cosine basis obtained by periodising a function on $[-1, 1]$ via the map $x = \cos \theta$ with $\theta \in [-\pi, \pi]$. Indeed, for any f on $[-1, 1]$ the composition $F(\theta) := f(\cos \theta)$ is an even 2π -periodic function, so it admits a cosine series $F(\theta) = \sum_{m \geq 0} a_m \cos(m\theta)$. Pulling this back to x gives a Chebyshev expansion in the polynomials

$$T_m(x) := \cos(m\theta) = \cos(m \arccos x), \quad x \in [-1, 1]. \quad (2.96)$$

They form an orthogonal system on $[-1, 1]$ with respect to the weight $w(x) = (1 - x^2)^{-1/2}$,

$$\int_{-1}^1 \frac{T_m(x) T_\ell(x)}{\sqrt{1 - x^2}} dx = \begin{cases} \pi, & m = \ell = 0, \\ \frac{\pi}{2}, & m = \ell \geq 1, \\ 0, & m \neq \ell, \end{cases} \quad (2.97)$$

which is the origin of the prefactor $1/(\pi\sqrt{1 - x^2})$ in the reconstruction formulas below. A key computational advantage is the three-term recurrence

$$T_0(x) = 1, \quad T_1(x) = x, \quad T_{m+1}(x) = 2x T_m(x) - T_{m-1}(x), \quad (2.98)$$

which lifts to operators by replacing $x \mapsto \hat{H}$. The Chebyshev polynomials are proven to be optimal in minimax approximation theory [144], making them well suited for spectral reconstructions.

For a target density $g(x)$ supported on $[-1, 1]$, define its Chebyshev moments

$$\mu_m = \int_{-1}^1 g(x) T_m(x) dx. \quad (2.99)$$

The truncated Chebyshev series reproduces $g(x)$ up to Gibbs oscillations. The KPM suppresses these oscillations by damping high orders with kernel coefficients $\gamma_m^{(M)}$,

yielding [136]

$$g_M(x) = \frac{1}{\pi\sqrt{1-x^2}} \left[\gamma_0^{(M)} \mu_0 + 2 \sum_{m=1}^M \gamma_m^{(M)} \mu_m T_m(x) \right]. \quad (2.100)$$

We use the Jackson kernel,

$$\gamma_m^{(M)} = \frac{(M-m+1) \cos\left(\frac{\pi m}{M+1}\right) + \sin\left(\frac{\pi m}{M+1}\right) \cot\left(\frac{\pi}{M+1}\right)}{M+1}, \quad (2.101)$$

which yields a positive approximation with controlled broadening [136]. Chebyshev methods require a bounded spectrum. With $E_{\max}, E_{\min}$ the extreme eigenvalues of the original Hamiltonian, define the rescaled operator

$$\hat{H}' = \frac{\hat{H} - a}{b}, \quad a = \frac{E_{\max} + E_{\min}}{2}, \quad b = \frac{E_{\max} - E_{\min}}{2 - \varepsilon}, \quad (2.102)$$

with a small $\varepsilon > 0$ to keep the spectrum away from ± 1 . All moments below refer to $\hat{H}'$; energies map back by $E = a + bx$. For notational simplicity we drop primes and write $\hat{H}$ for the rescaled operator. At fixed kernel, the truncation order M sets the effective energy resolution (after undoing the rescaling) as $\Delta E \approx \pi b/M$, up to kernel-dependent constants. For the (normalised) DOS $g(E) = \frac{1}{D} \sum_n \delta(E - E_n)$ of the rescaled $\hat{H}$, the moments are

$$\mu_m = \int_{-1}^1 g(E) T_m(E) dE = \frac{1}{D} \sum_n T_m(E_n) = \frac{1}{D} \text{Tr}[T_m(\hat{H})]. \quad (2.103)$$

Inserting these moments into Eq. (2.100) yields the Jackson-damped approximation $g_M(E)$. More generally, operator-resolved spectral weights follow from

$$\mu_{A,m} = \text{Tr}[\hat{A} T_m(\hat{H})], \quad (2.104)$$

which reconstruct the weighted density $h_A(E) = \text{Tr}[\hat{A} \delta(E - \hat{H})]$ as

$$h_{A,M}(E) = \frac{1}{\pi\sqrt{1-E^2}} \left[\gamma_0^M \mu_{A,0} + 2 \sum_{m=1}^M \gamma_m^M \mu_{A,m} T_m(E) \right]. \quad (2.105)$$

The special case $\hat{A} = |\psi\rangle\langle\psi|$ yields the LDOS of $|\psi\rangle$. Computing the direct traces in Eqs. (2.103) and (2.104) becomes infeasible at large D . Therefore, we use typicality, introduced in Sec. 2.9.3. For random probes $\{|r\rangle\}$ whose components satisfy the moment conditions Eq. (2.90),

$$\mu_m = \frac{1}{D} \text{Tr}[T_m(\hat{H})] \approx \frac{1}{R} \sum_{r=1}^R \langle r | T_m(\hat{H}) | r \rangle, \quad \mu_{A,m} \approx \frac{1}{R} \sum_{r=1}^R \langle r | \hat{A} T_m(\hat{H}) | r \rangle. \quad (2.106)$$

The matrix elements $\langle r | T_m(\hat{H}) | r \rangle$ are obtained without forming $T_m(\hat{H})$ explicitly by lifting the recursion Eq. (2.98) to vectors. For each probe $|r_0\rangle = |r\rangle$ define

$$|r_1\rangle = \hat{H} |r_0\rangle, \quad (2.107)$$

$$|r_m\rangle = 2\hat{H} |r_{m-1}\rangle - |r_{m-2}\rangle, \quad m \geq 2, \quad (2.108)$$

so that $\langle r | T_m(\hat{H}) | r \rangle = \langle r_0 | r_m \rangle$. The cost per probe is M sparse matrix-vector multiplications and $\mathcal{O}(M)$ vector operations, while memory scales as three working vectors $|r_{m-2}\rangle, |r_{m-1}\rangle, |r_m\rangle$.

In practical KPM calculations, M sets the nominal resolution $\Delta E \approx \pi b/M$ after undoing the rescaling, while kernel damping suppresses Gibbs oscillations at the price of broadening sharp features; accordingly, M is chosen to resolve the physical scale of interest rather than individual many-body levels. At fixed R , the remaining statistical uncertainty is governed by typicality and decreases rapidly with increasing Hilbert-space dimension (Sec. 2.9.3), so smooth densities often require only modest R , whereas resolving fine structure demands more probes. Numerical stability is controlled primarily by the rescaling margin ε , which keeps the spectrum away from ± 1 ; round-off accumulation at large M is usually mild for sparse matvecs but should be monitored when pushing to very deep recursions. Finally, the dominant classical cost is the repeated application of $\hat{H}$ (or its symmetry-restricted representation), so optimising the matvec primitive directly reduces wall-clock time.

2.9.5 Kernel polynomial method as a quantum algorithm




Figure 2.9: Circuit primitive used to estimate (approximated) Chebyshev moments on quantum hardware. A pseudo-random probe state is prepared on N system qubits and an ancilla is prepared in $|+\rangle$. A controlled unitary implementing $\hat{U}_{ST} = e^{\pm im\hat{H}_K}$ is applied, followed by single-qubit measurements on the ancilla to extract $\text{Tr}[\cos(m\hat{H}_K)]$ and $\text{Tr}[\sin(m\hat{H}_K)]$.

Classically, KPM reduces spectral reconstruction to estimating moments. On a quantum device the same reduction applies: the goal is to estimate moments expressed as traces (or averages) of functions of $\hat{H}$. The circuit in Fig. 2.9 implements a deterministic quantum computation with one clean qubit (DQC1)-like trace estimator for unitary evolutions whose real or imaginary parts approximate Chebyshev moments in the spectral bulk.




Figure 2.10: Pseudo-random circuit used to generate probe states on the system register. (a) Connectivity pattern for one instance at $N = 12$. (b) Gate decomposition into native entanglers and single-qubit rotations. The entangling layers use $\hat{R}^{zz}(\pi/2)$ gates, interleaved with local rotations chosen at random from $\{\hat{R}^x(\pi/2), \hat{R}^y(\pi/4), \hat{R}^z(\pi/2)\}$.

The DQC1-style estimator in Fig. 2.9 approximates the maximally mixed system register by sampling pseudo-random pure probes $|r\rangle$. We require an ensemble of states whose amplitudes satisfy the moment conditions in Eq. (2.90) sufficiently well, so that Eq. (2.92) is unbiased and exhibits typicality suppression at large Hilbert-space dimension. Rather than attempting Haar-random state preparation (which is depth-expensive), we generate *pseudo-random* probes by shallow random circuits, following the standard approach of alternating entangling layers with random single-qubit rotations (see Refs. [92, 145–147]).

For the Quantinuum H1-1 compilation model, we express the randomiser in terms of the native gates

$$\hat{R}^{zz}(\theta) = e^{-i\theta \hat{\sigma}^z \otimes \hat{\sigma}^z / 2}, \quad \hat{R}^z(\theta) = e^{-i\theta \hat{\sigma}^z / 2}, \quad (2.109)$$

and

$$\hat{U}_{1q}(\theta, \phi) = e^{-i\theta(\cos \phi \hat{\sigma}^x + \sin \phi \hat{\sigma}^y) / 2}. \quad (2.110)$$

We use $\hat{R}^{zz}(\pi/2)$ as the fixed entangling primitive and exploit all-to-all connectivity by choosing a sequence of pairings (“jumps”) across layers. One convenient choice is to connect sites $2i$ and $(2i + p_\ell) \bmod N$ with odd p_ℓ varying across layers, so that each qubit is rapidly scrambled without requiring deep circuits. Local $\hat{R}^z$ rotations can be implemented virtually by frame updates, and compilation can absorb commuting/anticommuting single-qubit gates into adjacent layers. We validate the scrambling quality of the randomiser by monitoring the half-system entanglement entropy. For $\hat{\rho}_r = |r\rangle\langle r|$ and a bipartition $A|B$ with $\dim \mathcal{H}_A = \dim \mathcal{H}_B$, we compute

$$s(\hat{\rho}_{r,A}) = -\text{Tr}[\hat{\rho}_{r,A} \ln \hat{\rho}_{r,A}], \quad \hat{\rho}_{r,A} = \text{Tr}_B[\hat{\rho}_r], \quad (2.111)$$

and compare it to the Page value $s_{\text{Page}} = \ln(\dim \mathcal{H}_A) - \frac{1}{2}$. In practice, we find that shallow depths already bring $s(\hat{\rho}_{r,A})$ close to s_{Page} over the system sizes considered, consistent with the probe ensemble behaving sufficiently “typical” for trace estimation.

In Fig. 2.11 we show representative data for system sizes $N = 10$ – 18 , averaged over

20 randomiser instances. On hardware, the stochastic-trace error due to imperfect typicality is often subdominant to finite-shot estimation of ancilla expectation values. Accordingly, increasing the number of circuit shots primarily reduces statistical uncertainty, while additional randomiser depth (or additional probe instances) offers diminishing returns once the probe ensemble has reached a sufficiently scrambled regime for the targeted observable.

In an ideal block-encoding or quantum singular value transformation (QSVT) setting, Chebyshev polynomials arise natively from a quantum-walk iterate and one can implement polynomial transforms of $\hat{H}$ with favourable asymptotic scaling in precision [148–150]. Here we keep the moment viewpoint but adopt a NISQ-oriented approximation that reduces moment estimation to traces of time-evolution operators. The full bridging argument is detailed in Sec. A.

Formally, if $\hat{U}_H = e^{i \arccos(\hat{H})}$, then

$$T_m(\hat{H}) = \cos(m \arccos(\hat{H})) = \frac{1}{2}(\hat{U}_H^m + \hat{U}_H^{-m}). \quad (2.112)$$

Implementing $\arccos(\hat{H})$ exactly is generally impractical, so we approximate it in the spectral bulk by truncating its odd-power expansion around 0,

$$\arccos(\hat{H}) = \frac{\pi}{2} - \sum_{k=0}^{\infty} c_k \hat{H}^{2k+1} \approx \frac{\pi}{2} - \hat{H}_K, \quad (2.113)$$

with $\hat{H}_K = \sum_{k=0}^K c_k \hat{H}^{2k+1}$. This yields, up to a known phase, traces of $e^{-im\hat{H}_K}$:

$$\cos(m(\frac{\pi}{2} - \hat{H}_K)) = \text{Re} \left[e^{im\pi/2} e^{-im\hat{H}_K} \right], \quad (2.114)$$

and therefore the proxy moments

$$\mu_m^K = \frac{1}{D} \text{Tr} \left[\cos(m(\frac{\pi}{2} - \hat{H}_K)) \right]. \quad (2.115)$$

At lowest order ($K = 0$), $\hat{H}_K = \hat{H}$ and the required primitive reduces to time evolution under $\hat{H}$.

The circuit of Fig. 2.9 is a DQC1-like trace estimator [151]: the ancilla implements a Hadamard test for $\text{Tr}[\hat{V}]$ with $\hat{V}$ a system unitary, while the system register approx-




Figure 2.11: (a) Convergence of half-chain von Neumann entropy to the Page value [84] for $N = 10$ – 18 , averaged over 20 random states. Shadows indicate standard deviation. (b) Comparison of relative entropy error ε_s at $N = 18$ for three circuit architectures: Richter *et al.* [147] (Ric.), parallel (Par., see Fig. 2.10), and sequential (Seq.) with respect to the number of layers of the random circuit.

imates the maximally mixed state by sampling pseudo-random pure states. In this way

$$\frac{1}{D} \text{Tr}[e^{-im\hat{H}_K}] \approx \frac{1}{R} \sum_{r=1}^R \langle r | e^{-im\hat{H}_K} | r \rangle, \quad (2.116)$$

and measuring the appropriate ancilla quadrature yields the real and imaginary parts, hence $\text{Tr}[\cos(m\hat{H}_K)]$ and $\text{Tr}[\sin(m\hat{H}_K)]$. Relative to classical stochastic trace estimation, the dominant additional uncertainty at fixed R is shot noise from estimating expectation values by repeated circuit executions.

To implement $e^{-im\hat{H}_K}$ we compile $\hat{H}_K$ into native gate layers and approximate time evolution using a Suzuki-Trotter (ST) product formula, which provides a transparent trade-off between depth and simulation error [152, 153]. Controlling the system evolution on the ancilla increases gate count and can raise effective locality (a k -local term becomes $(k+1)$ -local when controlled), motivating decompositions that minimise two-qubit interactions involving the control qubit (see Fig. 2.13).

For devices with limited connectivity, concentrating all controlled interactions onto a single ancilla can incur substantial SWAP overhead. A practical alternative

is to delocalise the control by preparing a GHZ state across several control qubits and using these as local controllers for disjoint neighbourhoods of the system register, as illustrated in Fig. 2.12. When $\hat{H}_K$ is geometrically local (e.g. nearest-neighbour), this layout can realise the controlled evolution with predominantly local two-qubit interactions and substantially reduces routing overhead.

Replacing the trace by $\text{Tr}[\hat{A}(\cdot)]$ yields operator-resolved moments, see Eq. (2.104). On hardware, one may either implement $\hat{A}$ coherently before the controlled simulation or apply a basis change after the controlled simulation and measure the qubits supporting $\hat{A}$ [154]. When the number of measured qubits increases, informationally complete measurements or classical shadows can reduce the sampling overhead [155, 156].




Figure 2.12: Delocalisation of the control for a chain of $N = 12$ qubits in a heavy-hex layout. The six orange qubits act as distributed controls for the N grey qubits. For nearest-neighbour Hamiltonians, each control qubit can govern two disjoint two-qubit interactions (yellow paths), avoiding SWAP overhead.

2.9.6 Implementation of the KPM on Quantinuum H1-1

This section reports the proof-of-principle implementation of the Chebyshev-moment estimation primitive on Quantinuum’s trapped-ion device (System Model H1-1) [157], where we used the non-integrable XYZ chain Eq. (2.37) as a hardware benchmark [1]. After the discussion in the previous sections on the quantum algorithm components, here we focus on (i) compilation choices specific to H1-1, and (ii) the empirical performance obtained with shallow approximations.

We benchmark the parameters $J_x = 1$, $J_y = 1/3$, $J_z = \Lambda = 1/2$. The staggered coupling breaks integrability without eliminating the exact symmetries of the model (in particular the $U(1)$ conservation of $\hat{S}^z = \sum_n \hat{s}_n^z$). In the hardware implementation it is convenient to compile controlled time-evolution (the controlled unitary in Fig. 2.9) using a first-order ST splitting into non-commuting layers acting on even-odd and odd-even bonds.




Figure 2.13: Gate decompositions used to implement the controlled unitary (i.e. the orange gate of Fig. 2.9) for the XYZ spin chain of Eq. (2.37). The controlled ST decomposition of this Hamiltonian requires a controlled $\exp(i\alpha\hat{s}_x \otimes \hat{s}^x + i\beta\hat{s}^y \otimes \hat{s}^y)$ for the odd-even spin pairs (a) and a controlled $\exp(i\alpha\hat{s}^x \otimes \hat{s}^x + i\beta\hat{s}^y \otimes \hat{s}^y + i\gamma\hat{s}^z \otimes \hat{s}^z)$ for

the even-odd spin pairs (b), with $\alpha, \beta, \gamma \in \mathbb{C}$. The hatched area in (a) represents the part that for the first ST step naturally mix with the last layer of the random circuit, without requiring additional gates; the hatched area of (b) in the last ST step can be neglected. The (c) and (d) respectively represent the conversion of the $C\text{-}\hat{R}^z(2\alpha)$ and $C\text{-}\hat{R}^x(2\alpha)$ in the native gate set currently supported by the System Model H1-1. Note that this particular gate decomposition was chosen in order to require as few gates as possible on the control qubit, to enhance parallelisation and reduce the probability of errors occurring.

A practical simplification arises from choosing $J_z = \Lambda$. Then

$$J_z + (-1)^n \Lambda = \begin{cases} 2J_z, & n \text{ even,} \\ 0, & n \text{ odd,} \end{cases} \quad (2.117)$$

so that one bond layer carries $xx + yy + zz$ interactions while the other carries only $xx + yy$ interactions. This reduces the controlled gate count without restoring the integrable structure of the uniform chain. The corresponding controlled two-qubit blocks are implemented with the decompositions shown in Fig. 2.13: odd-even pairs require a controlled $\exp(i\alpha\hat{\sigma}^x \otimes \hat{\sigma}^x + i\beta\hat{\sigma}^y \otimes \hat{\sigma}^y)$, while even-odd pairs require the additional zz rotation. As in Fig. 2.13, the compilation is chosen to minimise the number of gates acting on the control qubit, enabling greater parallelisation and reducing the exposure to control-line errors.




Figure 2.14: Estimation of the moments for $N = 12$ ((a)) and $N = 18$ ((b)) for the Hamiltonian of Eq. (2.37) with $J_x = 1$, $J_y = 1/3$, $J_z = \Lambda = 1/2$. The comparison is between Chebyshev moments obtained in different ways: analytically (Cheb), applying the arc-cosine approximation at the first order (Arccos), adding the ST approximation with one single step (ST), and the results from the circuit of Fig. 2.9 from Quantinuum, the emulator (H1-1e) and the quantum computer (H1-1). In (a) the estimates for the first 25 Chebyshev moments. We used 4 different random circuits ($R=4$) with $j = 5$ and 1000 shots. In (b) the estimates for the first 50 Chebyshev moments. Here, for the Quantinuum executions, we used 10 different random circuits ($R=10$) with $j = 4$ and 1000 shots. Each cross represents the result from a circuit with a different random circuit.

The random state used for trace estimation is generated by applying the parameterised randomiser of the previous Sec. 2.9.5 to the N -qubit register, and averaging over R independent random circuits. Chebyshev moments are then extracted as described already in the Sec. 2.9.5, with two deliberate truncations chosen to keep circuits shallow: (i) we approximate the arccosine mapping at order $K = 0$, i.e. $\arccos(\hat{H}) \approx \pi/2 - \hat{H}$, and (ii) we use a single ST step for the controlled evolution. Empirically, increasing K or the number of ST steps deepens the circuit but does not improve the final accuracy on hardware at fixed shot budget; in the present regime, gate errors dominate before the systematic approximation error becomes the bottleneck.

We executed the protocol for $N = 12$ and $N = 18$, corresponding to registers of 13 and 19 qubits respectively (including one control qubit) on H1-1. For each N we compare (i) exact moments from ED, (ii) moments obtained after the arccosine truncation, (iii) moments including the single-step ST approximation, and (iv) full circuit executions on the H1-1 emulator (H1-1e) and, where feasible, the physical device (H1-1); the data are shown in Figs. 2.14 and 2.15 and summarised at the level of DOS reconstructions in Fig. 2.1.

For $N = 12$ we estimate with different approximations the first $M = 25$ moments as presented in Fig. 2.14(a). The KPM reconstruction of the DOS appears already sufficient to match ED at the resolution of Fig. 2.15(a). Due to runtime constraints we use $R = 4$ random circuits, with 1000 shots per circuit, and resolve moments reliably




Figure 2.15: Reconstruction of the DOS for $N = 12$ (a) and $N = 18$ (b) for the Hamiltonian of Eq. (2.37) with $J_x = 1$, $J_y = 1/3$, $J_z = \Lambda = 1/2$. The different DOS are obtained using the five different estimates presented in Fig. 2.14 for the Chebyshev moments (using a kernel with $M = 25$ for $N = 12$ and $M = 50$ for $N = 18$) and the DOS from the ED. On both panels, on top: the DOS obtained with the five different estimates for the Chebyshev moments and the DOS from the ED. Since for $N = 18$ we computed μ_2 only from the H1-1 System, we also included what the DOS estimate from the μ_2 from the real hardware combined with the estimates from the emulator would look like (H1-1*). On bottom we reported the comparison of the DOS obtained with different approximation of the Chebyshev moments with respect to the analytical ones.

up to $m = 7$ on hardware. For $m > 7$ the moments are too small to be distinguished from zero with the available statistics; in the reconstruction we therefore set $\mu_m \simeq 0$ beyond this cutoff. Within this operating window, the hardware-extracted moments are consistent with the exact values and yield a DOS reconstruction that is accurate in the bulk of the spectrum.

For $N = 18$, accurate DOS reconstruction requires a higher expansion order, with $M \simeq 50$ in Fig. 2.14(b). We use $R = 10$ random circuits (again with 1000 shots each) and find that moments remain resolvable up to approximately $m = 11$ on the emulator at this shot budget. On the physical device, the increased circuit depth and cost at $N = 18$ restrict the feasible hardware runs; in practice we obtain only μ_2 directly from H1-1 and compute the remaining moments from the emulator. To illustrate the expected DOS in this mixed regime we report, in Fig. 2.15(b), a reconstruction that combines the hardware value of μ_2 with emulator estimates for higher orders.

Two empirical observations are worth recording. First, across both system sizes, H1-1e provides outcomes that are consistent with the exact (ED) benchmarks within the statistical uncertainty of the protocol, making it a reliable proxy for exploring compilation trade-offs at fixed depth. Second, the dominant two-qubit primitive in our compilation is the $\hat{R}^{zz}$ interaction (see Fig. 2.13), and the comparison between hardware and emulator suggests that the effective fidelity of this gate improves

for smaller rotation angles; residual discrepancies are consistent with an imperfect modelling of angle-dependent gate performance in the emulator.

Finally, all results in Figs. 2.14 and 2.15 are obtained without error mitigation. In this sense, the agreement visible in the bulk DOS should be read as a conservative proof of principle: with shallow truncations ($K = 0$ and one ST step) and modest sampling (shot-limited low-order moments), one can already access spectral information at system sizes beyond what is directly diagonalised in the hardware loop, and the resulting moment data integrate naturally into the spectral-density framework of this chapter.

2.9.7 Spectral reconstructions for discrete-time circuits

For Floquet dynamics the generator is unitary, so spectral information is naturally organised on the unit circle. Let $\hat{U}|\varphi_i\rangle = e^{-i\varphi_i}|\varphi_i\rangle$ with $\varphi_i \in [-\pi, \pi)$. In contrast to the Hamiltonian case, Chebyshev-Jackson KPM is not the natural basis for *global* spectral reconstructions: the orthogonal harmonics on the circle are Fourier modes. One can embed $\hat{U}$ into a Hermitian operator on $[-1, 1]$ via

$$\hat{A}_\phi = \frac{e^{i\phi}\hat{U} + e^{-i\phi}\hat{U}^\dagger}{2}, \quad (2.118)$$

whose eigenvalues are $a_i = \cos(\varphi_i - \phi)$, but the corresponding Chebyshev moments contain no information beyond the unitary trace moments,

$$\mu_m(\phi) = \frac{1}{D} \text{Tr}[T_m(\hat{A}_\phi)] = \frac{1}{D} \sum_i \cos(m(\varphi_i - \phi)) = \text{Re}(e^{im\phi} s_m), \quad s_m = \frac{1}{D} \text{Tr}[\hat{U}^m]. \quad (2.119)$$

Hence, for eigenphase densities and eigenphase-resolved correlators, Fourier reconstructions (with a suitable kernel on the circle) are the matched analogue of Chebyshev-Jackson expansions on $[-1, 1]$. Chebyshev methods become advantageous only when the target quantity is *itself* a spectral function of a Hermitian operator on $[-1, 1]$ (for instance, when one deliberately works with $\hat{A}_\phi$ as the object of interest, rather than using it as a detour back to eigenphases).

In Sec. 2.4.2 we discretise the staggered XYZ Hamiltonian Eq. (2.37) into a two-layer brickwork step

$$\hat{U}(\tau) = \prod_{n \text{ odd}} \hat{u}_{n,n+1}(\tau) \prod_{n \text{ even}} \hat{u}_{n,n+1}(\tau), \quad \hat{u}_{n,n+1}(\tau) = e^{-i\tau \hat{h}_{n,n+1}}, \quad (2.120)$$

with PBC (hence N even) and $\hat{h}_{n,n+1} = J_x \hat{s}_n^x \hat{s}_{n+1}^x + J_y \hat{s}_n^y \hat{s}_{n+1}^y + J_z^{(n)} \hat{s}_n^z \hat{s}_{n+1}^z$. For the data shown in Fig. 2.3 we take $N = 12$, $\tau = 1$, $J_x = 1$, $J_y = \frac{1}{3}$, and $(J_z, \Lambda) = (\frac{1}{2}, \frac{1}{2})$, so

that $J_{zz}^{\text{even}} = J_z + \Lambda = 1$ and $J_{zz}^{\text{odd}} = J_z - \Lambda = 0$.

The eigenphase DOS is

$$g(\varphi) = \frac{1}{D} \sum_{i=1}^D \delta(\varphi - \varphi_i), \quad D = 2^N, \quad (2.121)$$

with Fourier coefficients $s_m = \int_{-\pi}^{\pi} g(\varphi) e^{-im\varphi} d\varphi = \text{Tr}[\hat{U}^m]/D$. We reconstruct $g(\varphi)$ from the truncated, kernel-damped Fourier series

$$g_M(\varphi) = \frac{1}{2\pi} \left[1 + 2 \sum_{m=1}^M \gamma_m \text{Re}(s_m e^{im\varphi}) \right], \quad \gamma_m = \exp\left[-\frac{(m\sigma)^2}{2}\right], \quad (2.122)$$

which corresponds to smoothing by the circular heat kernel (a wrapped Gaussian) of angular width σ . In Fig. 2.3(a) we use $M = 80$ with $\sigma = 2\pi/M = \pi/40$, evaluate g_M on a uniform grid of 2048 points, and renormalise numerically so that $\int_{-\pi}^{\pi} d\varphi g_M(\varphi) = 1$. The ED reference is a density-normalised histogram of the eigenphases using 80 bins (a wrapped-Gaussian KDE with bandwidth 0.06 is included only as a visual guide).

For a local observable $\hat{A}$ and inverse temperature β (with respect to the Hamiltonian $\hat{H}$ entering the discretisation), we form the thermal Floquet autocorrelator moments

$$s_{A,m}^{(\beta)} = \frac{\text{Tr}[e^{-\beta\hat{H}} \hat{U}^{-m} \hat{A} \hat{U}^m \hat{A}]}{\text{Tr}[e^{-\beta\hat{H}}]}, \quad m \in \mathbb{Z}, \quad (2.123)$$

and reconstruct the corresponding eigenphase-difference density via the same heat-kernel Fourier series,

$$S_A^{(\beta)}(\varphi) \approx \frac{1}{2\pi} \left[s_{A,0}^{(\beta)} + 2 \sum_{m=1}^M \gamma_m \text{Re}(s_{A,m}^{(\beta)} e^{im\varphi}) \right], \quad \tilde{S}_A^{(\beta)}(\varphi) = \frac{S_A^{(\beta)}(\varphi)}{s_{A,0}^{(\beta)}}. \quad (2.124)$$

The normalisation ensures $\int_{-\pi}^{\pi} d\varphi \tilde{S}_A^{(\beta)}(\varphi) = 1$.

In Fig. 2.3(b) we take $\beta = 1$ and choose a representative *non-diagonal* two-site operator on sites (1, 2),

$$\hat{A} = \hat{\sigma}_1^x \hat{\sigma}_2^y - \hat{\sigma}_1^y \hat{\sigma}_2^x = 4(\hat{s}_1^x \hat{s}_2^y - \hat{s}_1^y \hat{s}_2^x), \quad (2.125)$$

embedded with identities on all other sites. The ED reference is computed by explicit summation over Floquet eigenphases with a 100-bin histogram (no additional smoothing). The typicality curve uses the thermal typicality construction introduced in Sec. 2.9.3: drawing R random probes $|r\rangle$ and defining $|\psi_r\rangle = e^{-\beta\hat{H}/2} |r\rangle$, we estimate

$$s_{A,m}^{(\beta)} \approx \frac{1}{R} \sum_{r=1}^R \frac{\langle \psi_r | \hat{U}^{-m} \hat{A} \hat{U}^m \hat{A} | \psi_r \rangle}{\langle \psi_r | \psi_r \rangle}, \quad (2.126)$$

with $R = 10$ probes and Fourier cutoff $M = 100$ (hence $\sigma = 2\pi/M$ in γ_m). For each probe one propagates $|\psi_r\rangle$ and $\hat{A}|\psi_r\rangle$ under repeated application of $\hat{U}$, accumulating the overlaps required by [Eq. \(2.126\)](#) for $m = 1, \dots, M$. The heat-kernel damping in [Eq. \(2.124\)](#) both regularises the truncated series and suppresses high- m statistical noise, which is the dominant source of uncertainty in typicality-based moment estimation.

Chapter 3

Finite-time dynamics

In [Ch. 2](#) we emphasised diagnostics defined as late-time limits: infinite-time averages, Drude weights, and spectral measures. Those objects are conceptually clean because they reduce long-time behaviour to spectral structure. They are, however, not the most natural outputs of a gate-based experiment or a digital simulator, which implements a finite number of Floquet cycles and returns stroboscopic time series of local observables. For this reason, the present chapter shifts viewpoint from asymptotic statements to *finite-time* transport diagnostics that are experimentally accessible in a $\hat{s}^z$ measurement basis.

This shift is also motivated by a broader perspective that was introduced in [Ch. 1](#) and will recur throughout the thesis: sufficiently large quantum circuits can be viewed as many-body systems in their own right, with emergent dynamical regimes that are not mere discretisations of Hamiltonian physics. In particular, discrete-time dynamics does not conserve energy, yet it can sustain other conservation laws (here magnetisation), so transport becomes one of the cleanest probes of “digital phases of matter” and of their crossovers.

The goal is twofold. First, we want practical probes that distinguish localised, diffusive, and anomalous dynamical regimes using only quantities that are straightforward to measure in a $\hat{s}^z$ basis. Second, we want to expose intermediate-time structure that is invisible to infinite-time limits. In the present model this is essential: near the integrable SWAP point the dynamics develops a long-lived window in which excitations propagate coherently and quasi-ballistically before eventual equilibration at much later times.

We develop a local-quench protocol tailored to the U(1)-symmetric Floquet circuits introduced in [Sec. 2.6](#). We initialise a single localised magnetisation inhomogeneity on an otherwise high-temperature background and follow the resulting magnetisation profile under repeated application of the Floquet unitary. While the raw profile has a transparent hydrodynamic interpretation, it is not an efficient object from which to extract robust scaling diagnostics under disorder averaging and periodic boundary conditions. The main technical step of the chapter is therefore to compress the full profile into the lowest circular moment of an associated (quasi-)probability distribution. This reduction yields two scalar time series with direct transport meaning: a drift variable (the phase velocity of the slowest Fourier mode) and a spread variable (its decay

rate), providing a compact and numerically stable summary of finite-time transport.

The chapter mirrors the structure of [Ch. 2](#): we first introduce finite-time concepts and diagnostics, then present benchmarks in the same discrete-time XXZ circuit models introduced in [Secs. 2.5, 2.6 and 2.7.2](#), and finally summarise the numerical tools used to produce the data. Concretely, [Sec. 3.1](#) motivates prethermalisation as a controlled finite-time window; [Sec. 3.2](#) defines the local-quench protocol and its reduction to circular moments; [Secs. 3.3 to 3.6](#) extract transport summaries across the phase diagram and quantify the near-SWAP “swappy” regime; and [Sec. 3.7](#) revisits the single-impurity benchmark of [Sec. 2.7](#) from a quasi-conserved-charge perspective. For comparison with the spectral and linear-response viewpoint of [Ch. 2](#), we occasionally comment on how the same finite-time information is encoded in current autocorrelations and in the low-frequency structure of the regular conductivity. We close with a discussion in [Sec. 3.8](#) and collect method details in [Secs. 3.9 and 3.9.2](#).

The results presented in this chapter build on Ref. [\[1\]](#), produced during this PhD, and on ongoing work [\[5\]](#). The local-quench protocol and the circular-moment transport diagnostics ([Sec. 3.2](#)), as well as the characterisation of the near-SWAP prethermal (“swappy”) regime ([Sec. 3.4](#)), follow Ref. [\[1\]](#). In addition, this chapter introduces and analyses a discrete-time single-impurity XXZ circuit prethermal plateau, formulated in terms of quasi-conserved charges and defect-controlled leakage ([Sec. 3.7](#)). These results will be presented in Ref. [\[5\]](#).

3.1 Floquet prethermalisation as a finite-time window

In periodically driven many-body systems, long-time stroboscopic dynamics within each exact symmetry sector is generically expected to approach the infinite-temperature state. By *Floquet prethermalisation* we mean a controlled separation of time scales: local observables relax on a microscopic time t_{loc} to a quasi-stationary regime, and only on a parametrically longer time t_* do they drift away towards infinite temperature. Crucially, the plateau is governed by an *emergent quasi-conserved local structure* (often an extensive generator, but not necessarily), and the lifetime t_* diverges in a controlled limit. Throughout, we use “prethermal” in this strict sense: a window $t_{\text{loc}} \ll t \ll t_*$ whose longevity is controlled by a quasi-conserved local structure, not merely by finite-size recurrences.

Two distinct mechanisms produce such prethermal windows in Floquet settings. One is the high-frequency mechanism, where t_* is exponentially large in the drive frequency and the plateau is organised by a dressed effective Hamiltonian $\hat{H}_*$. Another is the

near-integrable mechanism, where t_* diverges as a power of a detuning from an exactly solvable Floquet point; this is the sense relevant to the near-SWAP “swappy” regime studied in [Sec. 3.4](#).

3.1.1 High-frequency prethermalisation and a dressed generator $\hat{H}_*$

For broad classes of local lattice drives with bounded microscopic scales, rigorous results show that at sufficiently high drive frequency there exists an extensive, quasi-local operator $\hat{H}_*$ such that the Floquet evolution approximately conserves it over exponentially long times [[34](#), [35](#), [158–160](#)]. Equivalently, one may formulate this as an exponentially small dressed commutator and an exponentially long heating time,

$$\|[\hat{U}, \hat{H}_*]\| \sim e^{-c\omega}, \quad t_* \sim e^{+c\omega}, \quad (3.1)$$

with ω the drive frequency. On the window $t_{\text{loc}} \ll t \ll t_*$, local physics is well approximated by evolution under $\hat{H}_*$, and quasi-stationary values of local observables are captured by a Gibbs-like state of $\hat{H}_*$ at an effective inverse temperature fixed by the approximately conserved $\langle \hat{H}_* \rangle$.

A practical consequence is that a chosen “reference” Hamiltonian $\hat{H}$ (e.g. the continuous-time model one heuristically associates to the circuit) is not conceptually protected as a heating diagnostic: prethermal plateaux are defined by quasi-conservation of a *local* $\hat{H}_*$, not by stationarity of an arbitrarily selected $\hat{H}$ [[35](#), [160](#)]. Accordingly, when we discuss high-frequency prethermal behaviour we phrase plateaux in terms of $\hat{H}_*$ (or local densities that fix an effective temperature with respect to $\hat{H}_*$), rather than in terms of a specific microscopic $\hat{H}$.

3.1.2 Prethermalisation near an integrable Floquet point

A second, conceptually independent route to a long-lived intermediate-time regime is proximity to an exactly solvable Floquet dynamics. This mechanism does not rely on a high-frequency limit and applies equally to continuous- and discrete-time settings: one starts from an integrable reference evolution and considers a weak perturbation that breaks the integrable structure. The resulting dynamics can exhibit a pronounced separation of time scales: a rapid approach to a quasi-stationary state constrained by (approximately) conserved integrable structures, followed only at much later times by crossover to the generic long-time behaviour [[161–163](#)].

To state this in a model-agnostic way, let $\hat{U}_0$ denote a reference Floquet step for which

the dynamics is integrable in the sense that there exists an extensive set of algebraically independent (quasi-)local conserved charges $\{\hat{Q}_a\}$ sufficient to enforce GGE-like local constraints rather than direct relaxation to the symmetry-only infinite-temperature state. Consider a one-parameter deformation $\hat{U}(\lambda)$ with $\hat{U}(0) = \hat{U}_0$, where λ controls the strength of an integrability-breaking perturbation. On general grounds, integrability breaking dresses the conserved quantities into quasi-conserved operators $\hat{Q}_{*,a}(\lambda)$ whose stroboscopic drift per period is parametrically small as $\lambda \rightarrow 0$. A convenient diagnostic is therefore the slow evolution of the constrained data $\langle \hat{Q}_{*,a}(\lambda; t) \rangle$ compared with the microscopic relaxation time t_{loc} .

The hallmark of near-integrable prethermalisation is the divergence of a crossover time scale $t_*(\lambda)$ as $\lambda \rightarrow 0$. Unlike high-frequency Floquet prethermalisation, where t_* can be exponentially large in the drive frequency and the plateau is organised by a dressed effective Hamiltonian, here t_* is typically controlled algebraically by the integrability-breaking parameter. In the generic situation in which the perturbation induces scattering between integrable sectors at a Golden-Rule rate, one expects

$$\Gamma(\lambda) \propto \lambda^2, \quad t_*(\lambda) \sim \Gamma(\lambda)^{-1} \propto \lambda^{-2}, \quad (3.2)$$

up to model-dependent prefactors and possible modifications from resonances, selection rules, or additional symmetries. The intermediate-time regime may be characterised either by the near-stationarity of a suitably chosen set of quasi-conserved quantities, or by the quasi-stationarity of a broad class of local observables compared to their eventual long-time values, with both descriptions becoming asymptotically equivalent when the integrable constraints are sufficiently complete.

In this thesis we use the term *prethermal* in this near-integrable sense whenever an extended time window is controlled by weak breaking of an exactly solvable Floquet structure and the associated crossover time diverges in a controlled limit.

3.2 Transport diagnostics from magnetisation profile

As in this chapter we focus on extracting finite-time transport information from the spreading of a local $U(1)$ charge inhomogeneity under a $U(1)$ -symmetric Floquet unitary $\hat{U}$, we first define the relevant observables and then explain how to reduce them to compact transport diagnostics. The prethermal “swappy” regime discussed below is an instance of the near-integrable mechanism described in [Sec. 3.1.2](#). The organising viewpoint is the same as in the Hamiltonian transport literature: transport is encoded either (i) in equilibrium current correlators (Green-Kubo) or, equivalently,

(ii) in the spatiotemporal spreading of the conserved density (dynamic structure factors and their moments) [43]. The present section formalises the second route in a form tailored to discrete-time circuits on a ring.

3.2.1 Local magnetisation profile in a fixed- M sector

We work stroboscopically with $t \in \mathbb{Z}$ the Floquet-cycle index and define $\hat{s}_n^z(t) = \hat{U}^{-t} \hat{s}_n^z \hat{U}^t$. Given an initial state $\hat{\rho}(0)$ supported on the fixed-magnetisation subspace $\mathcal{H}^{(M)}$, the local magnetisation profile is

$$M_n(t) = \text{Tr} [\hat{\rho}(0) \hat{s}_n^z(t)], \quad n \in \{1, \dots, N\}. \quad (3.3)$$

Charge conservation implies $\sum_n M_n(t) = M$ for all t .

The protocol below uses a localised charge perturbation on top of an infinite-temperature background within $\mathcal{H}^{(M)}$. Let $\hat{P}_M$ project onto $\mathcal{H}^{(M)}$ and let $\hat{P}_{n_0}^\dagger = \hat{I}/2 + \hat{s}_{n_0}^z$ polarise a chosen site n_0 (we take $n_0 = N/2$). A convenient mixed initial condition is

$$\hat{\rho}(0) = \frac{\hat{\tilde{\rho}}}{\text{Tr}[\hat{\tilde{\rho}}]}, \quad \hat{\tilde{\rho}} = \hat{P}_M \left[\hat{P}_{n_0}^\dagger \otimes \left(\frac{\hat{I}_2}{2} \right)^{\otimes(N-1)} \right] \hat{P}_M. \quad (3.4)$$

Eq. (3.4) is the sector-restricted analogue of the standard “local quench at infinite temperature” used to measure the spreading of equilibrium correlations; see Sec. C for the explicit connection to sector-resolved correlators and typicality.

3.2.2 Semiprobability profile in the $M=0$ sector

To reduce $M_n(t)$ to compact transport diagnostics, it is useful to normalise it into a unit-mass profile. From here on we specialise to the largest sector $M=0$. For the initial condition Eq. (3.4) with one fully polarised spin, the $M=0$ constraint fixes the (uniform) background magnetisation on $n \neq n_0$ as

$$M_{n_0}(0) = \frac{1}{2}, \quad M_{n \neq n_0}(0) = M_B = -\frac{1}{2(N-1)}. \quad (3.5)$$

We then define the semiprobability profile

$$p_n(t) = 2 \frac{M_n(t) - M_B}{1 - 2M_B}, \quad \sum_{n=1}^N p_n(t) = 1, \quad (3.6)$$

so that $p_{n_0}(0)=1$ and $p_{n \neq n_0}(0)=0$.

In the transport literature, the object most commonly analysed is a normalised density-density correlator (or structure factor) whose zeroth moment is fixed to unity [43, 113, 164]. Eq. (3.6) is precisely this normalisation, specialised to a sector-restricted local-quench state. At finite N , typicality fluctuations can yield $M_n(t) < M_B$ for some sites, making $p_n(t)$ weakly negative. In that case $p_n(t)$ should be interpreted as a semiprobability; this does not obstruct moment-based diagnostics and becomes negligible as $\dim \mathcal{H}^{(0)}$ grows.

3.2.3 From spreading profiles to hydrodynamic transport

The relevance of $p_n(t)$ is that, in the hydrodynamic regime, its long-wavelength evolution is constrained. In the simplest case of unbiased diffusion on a line one expects the scaling form

$$p_n(t) \simeq t^{-1/z} f\left(\frac{n - n_0}{t^{1/z}}\right), \quad (3.7)$$

with dynamical exponent z (diffusion: $z = 2$, ballistic: $z = 1$; KPZ superdiffusion corresponds to $z = 3/2$ in the isotropic Heisenberg chain) [43, 109, 110]. Two standard consequences are: (i) a width $\sigma(t) \sim t^{1/z}$ extracted from second moments, and (ii) a return/peak scale $p_{\max}(t) \sim t^{-1/z}$ for unimodal profiles. This motivates the exponents introduced below.

A complementary, Green-Kubo perspective is obtained by combining the continuity equation with linear response. In continuous time one may relate the growth of the mean-square displacement of the density correlator to the time integral of the current autocorrelator (Einstein/Green-Kubo relations) [43]. In discrete time the same logic applies with finite differences and a stroboscopic continuity equation; we use the profile viewpoint because it remains numerically stable in strongly disordered Floquet circuits and because it is directly comparable to experimentally accessible density profiles.

3.2.4 Circular ensemble and its relation to the dynamic structure factor

With periodic boundary conditions, linear moments become ambiguous once the profile approaches wrap-around. This is especially relevant near SWAP-like points, where an excitation can propagate almost ballistically and remain coherent up to times comparable with the wrap-around time t_{wrap} (defined in Eq. (3.16)). A packet that is still well localised on the ring may then appear, in a linear site coordinate, as two peaks separated by the arbitrary coordinate cut (near $n = 1$ and $n = N$), so that the apparent centre-of-mass and variance become gauge-dependent artefacts. We therefore encode $p_n(t)$ by its lowest nonzero Fourier component (a standard object in transport

theory: the dynamic structure factor at small momentum) [43]. Define

$$R_k(t) = \sum_{n=1}^N p_n(t) e^{ik(n-n_0)}, \quad k = \frac{2\pi m}{N}, \quad m \in \mathbb{Z}, \quad (3.8)$$

and focus on the smallest momentum $k_1 = 2\pi/N$. Our first circular moment is exactly

$$R(t) = R_{k_1}(t) = \sum_{n=1}^N p_n(t) e^{i\theta_n}, \quad \theta_n = \frac{2\pi}{N}(n - n_0). \quad (3.9)$$

In a diffusive system, hydrodynamics predicts

$$R_k(t) \simeq e^{-Dk^2 t} \quad (3.10)$$

for small k , where D is the diffusion constant. If, in addition, the profile is *advected* with a finite velocity v (sometimes called *drift velocity*, *bias*, or *convective/advection velocity* depending on context), then the long-wavelength mode obeys an advection-diffusion form

$$R_k(t) \simeq e^{ivkt} e^{-Dk^2 t}. \quad (3.11)$$

Eq. (3.11) is the precise sense in which our “drift” diagnostic is a transport quantity: it is the phase velocity of the slowest Fourier mode of the spreading packet, i.e. the convective component of the hydrodynamic mode.

We extract from $R(t)$ a circular location and a circular spread,

$$\mu(t) = \text{Arg } R(t), \quad \sigma(t) = \sqrt{-2 \log |R(t)|} \geq 0, \quad (3.12)$$

which is an exact parametrisation of the complex number $R(t)$. When Eq. (3.11) holds at $k=k_1$, it implies

$$\mu(t) \simeq vk_1 t, \quad \sigma(t)^2 \simeq 2Dk_1^2 t. \quad (3.13)$$

Thus, on time windows where the lowest mode is hydrodynamic, one can extract *effective* transport coefficients directly:

$$v_{\text{eff}}(t) = \frac{1}{k_1} \text{Arg}(R(t+1)R(t)^*), \quad D_{\text{eff}}(t) = -\frac{1}{k_1^2} \log \left| \frac{R(t+1)}{R(t)} \right|. \quad (3.14)$$

In our notation below we report the magnitude of the one-step phase increment as

$$\nu(t) = \frac{N}{2\pi} \left| \text{Arg}(R(t+1)R(t)^*) \right| = |v_{\text{eff}}(t)|, \quad (3.15)$$

which measures displacement in sites per Floquet cycle. The product $R(t+1)R(t)^*$ removes branch-cut artefacts that appear if one naively differences $\mu(t)$ across $\pm\pi$. With this normalisation, the brickwork light-cone bound corresponds to $\nu(t) \lesssim 2$.

Because the geometry is a ring, sufficiently fast coherent drift inevitably leads to finite-size wrap-around and recurrences. We define a characteristic wrap-around time as the first time at which the cumulative drift distance reaches the system size,

$$t_{\text{wrap}} := \inf \left\{ t \in \mathbb{N} \left| \sum_{s=0}^{t-1} \nu(s) \geq N \right. \right\} \approx \frac{N}{\nu_{\text{typ}}}, \quad (3.16)$$

where ν_{typ} is a representative drift speed on the time window of interest.

On large rings (small k_1) the cumulants of the packet can be recovered from $\log R_k(t)$ near $k=0$: the first cumulant is the centre-of-mass displacement and the second cumulant is the variance. Our construction samples this cumulant-generating function at $k=k_1$. This is why $\mu(t)$ and $\sigma(t)$ are robust substitutes for the linear mean and width once wrap-around becomes relevant.

3.2.5 Wrapped-normal reference and effective exponents

In the diffusion equation on a line, an initially localised profile broadens as a Gaussian with width $\propto t^{1/2}$. On a ring, the natural analogue is a wrapped Gaussian,

$$\mathcal{N}_{\text{W}}(n; \mu, \sigma) \propto \sum_{k \in \mathbb{Z}} \exp \left[-\frac{(\theta_n - \mu + 2\pi k)^2}{2\sigma^2} \right], \quad (3.17)$$

whose first circular moment satisfies $R = e^{i\mu - \sigma^2/2}$, matching [Eq. \(3.12\)](#). This motivates [Eq. \(3.12\)](#) as a spread diagnostic when $p_n(t)$ remains approximately unimodal. A limitation is that $\sigma(t)$ tracks a physical width reliably only in that regime. In the later analysis we identify a “swappy” window in which the excitation remains coherent and quasi-ballistic up to times comparable with the wrap-around time t_{wrap} ([Eq. \(3.16\)](#)), so that width-based diagnostics must be interpreted with care. In this regime the profile can develop multi-peak and staggered structures at early times. Accordingly, the most robust transport signature becomes the convective component $\nu(t)$ (equivalently v_{eff}).

To quantify broadening and decay we also track the peak height $p_{\text{max}}(t) = \max_n p_n(t)$ and extract effective exponents on intermediate-time windows,

$$\sigma(t) \sim t^{\alpha_\sigma}, \quad p_{\text{max}}(t) \sim t^{-\alpha_p}. \quad (3.18)$$

They are finite-time, finite-size versions of [Eq. \(3.7\)](#). For unimodal packets one expects

$$\alpha_\sigma \simeq \alpha_p \simeq \frac{1}{z}.$$

Thus $1/z = 1/2, 1, 2/3$ for diffusion, ballistic spreading, and KPZ scaling, respectively.

In particular, α_p is closely related to the decay of local autocorrelations at the origin (return probability), which is frequently used as a transport diagnostic when extracting full profiles is costly [43, 164].

3.3 Disordered Floquet XXZ circuit




Figure 3.1: Dynamics of a single spin inhomogeneity across representative regimes with $J_z=\pi$ and at particular values of J , symbols as defined in Sec. 2.6.1. Results of numerical propagation of the initial state Eq. (3.4) (implemented via typicality as in Sec. 3.9.2) in the $M=0$ sector under the disordered Floquet XXZ circuit Eq. (3.19). Top row: typical single realisations. Bottom row: disorder-averaged profiles (100 realisations). Solid/dashed/dotted horizontal lines mark $t=1$, $t=N/2$ (ballistic wrap-around benchmark; Eq. (3.16)), and $t=(N/2)^2$ (diffusive wrap-around estimate), respectively.

Throughout the next sections we study the disordered discrete-time XXZ chain introduced in Sec. 2.6. We briefly recall its Floquet construction to fix notation (in particular the SWAP point at $J=\pi$ discussed in Sec. 2.5).

Across this phase diagram, the circular-moment diagnostics introduced above yield transport exponents consistent, at accessible sizes and times, with phenomenological localisation, ordinary diffusion, and an anomalous superdiffusive window near the SWAP point; the latter coincides with the near-integrable prethermal “swappy” regime discussed in Sec. 3.4.

One Floquet cycle is generated by a $U(1)$ -symmetric unitary $\hat{U}$ built from nearest-neighbour gates,

$$\hat{U} = \prod_{\{P(n)\}} \hat{u}_{n,n+1}, \quad (3.19)$$

where the product runs over a random permutation P of the bonds $(n, n+1)$ with

periodic boundary conditions. The two-site gate is

$$\hat{u}_{n,n+1} = e^{-i\hat{h}_{n,n+1}} e^{-i\hat{h}_n \hat{s}_n^z} e^{-i\hat{h}'_n \hat{s}_{n+1}^z}, \quad (3.20)$$

with

$$\hat{h}_{n,n+1} = \frac{J}{2} (\hat{s}_n^+ \hat{s}_{n+1}^- e^{i\phi_n} + \hat{s}_n^- \hat{s}_{n+1}^+ e^{-i\phi_n}) + J_z \hat{s}_n^z \hat{s}_{n+1}^z. \quad (3.21)$$

The parameters $\{h_n, h'_n, \phi_n\}$ are sampled independently and uniformly from $(-\pi, \pi]$ for each gate, while J and J_z are fixed across all gates in a given Floquet unitary. By construction $[\hat{U}, \hat{S}^z]=0$, so we work in fixed- M sectors and focus on $M=0$ throughout.

The local magnetisation profile Eq. (3.3) is the most direct finite-time observable in this model: it is local, stroboscopic, and experimentally accessible. Figure 3.1 shows representative realisations and disorder-averaged profiles at $N=20$ for four points along the $J_z=\pi$ line. The horizontal lines indicate: one cycle ($t=1$), a ballistic wrap-around benchmark ($t=N/2$, cf. t_{wrap} in Eq. (3.16)), and a diffusive wrap-around estimate ($t=(N/2)^2$), which serve as heuristic benchmarks before we quantify drift and broadening via circular moments.

Time evolution is carried out by repeated application of $\hat{U}$ to a state vector. We evaluate each of the first $t=100$ cycles explicitly (cycle-by-cycle sampling). To reach later times, we optionally apply blocks of Floquet cycles and sample the dynamics more coarsely for $t>100$. This can create an apparent late-time “pinning” in plots if one forgets the sampling schedule, but it is purely numerical: the physical dynamics is always generated by single-cycle applications of $\hat{U}$. The subsequent analysis is formulated in terms of stroboscopic time t and is insensitive to the sampling strategy once fit windows are chosen within well-resolved intervals.

3.4 Prethermalisation near the SWAP point: the swappy regime




Figure 3.2: Prethermal scaling near the SWAP point. Values are extracted from $\langle \sigma(t) \rangle$ as $t_\sigma = \inf\{t \mid \langle \sigma(t) \rangle > N/30 + 1.25\}$ (panel (a)), and from $\langle p_{\text{max}}(t) \rangle$ as $t_p = \inf\{t \mid \langle p_{\text{max}}(t) \rangle < 2.6/N\}$ (panel (b)). Dotted lines show power-law fits $t_\sigma \propto (J')^{-2.4}$ and $t_p \propto (J')^{-2.6}$.

Here “prethermal” is used in the near-integrable sense of [Sec. 3.1.2](#), controlled by detuning $J'=\pi-J$ from the integrable SWAP circuit. This is distinct from high-frequency $\hat{H}_*$ prethermalisation of [Sec. 3.1.1](#).

At the SWAP point $J=\pi$, a local $\hat{s}^z$ inhomogeneity is transported coherently by near-perfect exchange of neighbouring excitations. In particular, the first circular moment $R(t)$ undergoes an approximately rigid rotation with $|R(t)|\simeq 1$, so that the circular variance $\sigma(t)=\sqrt{-2\log|R(t)|}$ remains small, while the peak height $p_{\max}(t)$ stays large. Detuning away from $J=\pi$ introduces weak integrability breaking: SWAP-like ballistic motion persists at early and intermediate times, but rare scattering/dephasing events gradually convert coherent advection into irreversible broadening and peak decay.

The role of the timescales t_p and t_σ is to quantify, in a finite-size ring geometry, the *breakdown time* of SWAP-like coherent transport using two complementary observables. First, $p_{\max}(t)$ is a local coherence-sensitive diagnostic: it tracks the survival of a sharp packet and is closely related to the return probability (local autocorrelation) at the origin. Second, $\sigma(t)$ is a long-wavelength diagnostic: since $\sigma(t)$ is extracted from the magnitude of the lowest Fourier component $R_{k_1}(t)$, it measures the decay of the slowest mode and remains numerically stable even when the real-space profile develops staggered or multi-peak structure (where any linear “width” becomes gauge-dependent on a ring). We therefore define two threshold-based crossover times,

$$t_p = \inf\{t \in \mathbb{N} \mid \langle p_{\max}(t) \rangle < p_{\text{th}}\}, \quad (3.22)$$

$$t_\sigma = \inf\{t \in \mathbb{N} \mid \langle \sigma(t) \rangle > \sigma_{\text{th}}\}, \quad (3.23)$$

and interpret them as indicators for the onset of irreversible broadening (loss of SWAP-like coherence), not as heating-to-infinite-temperature times. Geometric caveat: thresholds and fit windows are chosen with $t \ll t_{\text{wrap}}$ (see [Eq. \(3.16\)](#)) so that wrap-around recurrences on the ring do not masquerade as long-lived plateaux.

The near-integrable expectation for the divergence of these crossover times follows from a perturbative rate argument. Writing the detuned circuit as $\hat{U}=\hat{U}_0 e^{-iJ'\hat{V}}$ and moving to the interaction picture with respect to $\hat{U}_0$, the detuning produces scattering of the integrable SWAP-like motion. Transition amplitudes are linear in J' , hence transition probabilities per cycle are generically quadratic. Accordingly, one expects a relaxation rate $\Gamma(J') \propto (J')^2$ and a breakdown time

$$t_*(J') \sim \Gamma(J')^{-1} \propto (J')^{-2}, \quad (3.24)$$

up to model-dependent prefactors and finite-size corrections. Both t_p and t_σ probe this same underlying rate: t_p through the survival of a sharp local feature, and t_σ through the decay of the lowest Fourier component $|R(t)|$ (equivalently the growth of

$$\sigma(t)^2 = -2 \log |R(t)|.$$

The data in Fig. 3.2 show clear power-law growth of both crossover times as $J' \rightarrow 0$. The fitted exponents depend mildly on the threshold choice and on the finite-size/time window, but across N and across reasonable thresholds we consistently find a divergence compatible with the Golden-rule baseline $t \propto (J')^{-2}$. Thus the swappy regime is a genuinely dynamical ergodic-swappy crossover (seen in the finite-time observables of Ch. 3 near $J \lesssim \pi$ but absent from the static analyses of Sec. 2.1) with near-integrable prethermal scaling $t_* \propto (\pi - J)^{-2}$ as $J \rightarrow \pi$ (Sec. 3.4 and Fig. 3.2) [162, 165].

3.5 Transport observables from drift, spreading, and peak decay

From each trajectory $R(t)$ we derive the spread $\sigma(t)$ and drift speed $\nu(t)$ via Eqs. (3.12) and (3.15). To complement $\sigma(t)$ we also track the instantaneous peak height $p_{\max}(t) = \max_n p_n(t)$. These three observables distinguish drift-dominated dynamics from diffusive broadening in a compact and statistically stable way.

To quantify spreading and relaxation we extract effective power laws on intermediate-time windows,

$$\sigma(t) \sim t^{\alpha_\sigma}, \quad p_{\max}(t) \sim t^{-\alpha_p}, \quad (3.25)$$

with diffusion benchmark $\alpha_\sigma = \alpha_p = 1/2$ and ballistic benchmark $\alpha_\sigma = \alpha_p = 1$.




Figure 3.3: Realisation-averaged dynamics of transport observables along the $J_z = \pi$ line. Shown are (a) the spread $\langle \sigma(t) \rangle$, (b) the drift speed $\langle \nu(t) \rangle$, and (c) the peak height $\langle p_{\max}(t) \rangle$ for various $J \in (0, \pi)$ at fixed $J_z = \pi$. Grey-shaded regions indicate the intermediate-time windows used for extracting α_σ and α_p . Due to stroboscopic sampling limitations, drift is displayed only up to the time range where cycle-by-cycle data are available; this captures the relevant transient behaviour.

The spread and peak decay provide robust diffusion diagnostics in the ergodic regime, while the swappy and near-SWAP regimes are cleanly identified by persistent drift at speeds of order one site per cycle.

3.6 Extracting effective exponents and drift velocities

We now process the realisation-averaged observables $\langle\sigma(t)\rangle$, $\langle\nu(t)\rangle$,^(a) and $\langle p_{\max}(t)\rangle$ to obtain finite-size diagnostics across $N \in \{14, 16, 18, 20, 22\}$. To determine α_σ from $\langle\sigma(t)\rangle \sim t^{\alpha_\sigma}$ we fit on intermediate-time windows that exhibit stable power law behaviour (shaded in Fig. 3.3(a)).¹

The extracted α_σ displays two robust plateaus within accessible sizes: $\alpha_\sigma \simeq 0$ (suppressed spreading) at small J and $\alpha_\sigma \simeq 1/2$ (diffusive broadening) over an extended range at larger J . As J enters the swappy window, α_σ becomes unstable and error bars grow, as expected from the non-unimodal structure of $p_n(t)$ in that regime; drift-based diagnostics remain stable precisely in the same parameter window. Close to the SWAP point $J \rightarrow \pi$, α_σ rapidly returns towards 0, consistent with strongly suppressed broadening. To summarise coherent transport we define a time-averaged drift speed




Figure 3.4: Extracted transport summaries along the $J_z=\pi$ line. (a) Effective spread exponent α_σ extracted from $\langle\sigma(t)\rangle$. (b) Time-averaged drift speed $\langle\nu\rangle$ defined in Eq. (3.26).

$$\langle\nu\rangle = \frac{1}{T} \sum_{t=0}^{T-1} \langle\nu(t)\rangle, \quad (3.26)$$

where T is chosen as the time for a typical near-SWAP trajectory to return to its initial position on the ring. Since the light-cone bound is $\nu(t) \lesssim 2$ sites per cycle, a natural choice is

$$T = t_{\text{wrap}} \approx \frac{N}{\nu_{\text{typ}}}, \quad \nu_{\text{typ}} \simeq 2, \quad (3.27)$$

¹To reduce sensitivity to the choice of fit window and to disorder fluctuations, we perform an ensemble of fits on slightly perturbed time intervals. Concretely, for each of 25 runs we (i) perturb the endpoints of the nominal fit window by independent normally distributed factors with mean 1 and standard deviation 0.2, (ii) compute $\langle\sigma(t)\rangle$ using a random one-third subset of trajectories, and (iii) fit $\log\langle\sigma(t)\rangle = b \log t + c$. We then report α_σ as the mean of the resulting slopes b , with the standard deviation as an error estimate. The same procedure is applied to $p_{\max}(t)$ to estimate α_p .

rounded to the nearest integer in numerics. The resulting $\overline{\langle \nu \rangle}$ is reported in Fig. 3.4(b).

We also extract the peak-decay exponent α_p from $\langle p_{\max}(t) \rangle \sim t^{-\alpha_p}$ using the same fitting protocol. Away from the $J_z = \pi$ line, some parameter values yield effective $\alpha_p \simeq 1$ or larger (very fast relaxation of local observables); the interpretation and limitations of such large effective exponents are discussed together with the prethermal analysis in Sec. 3.4.

3.7 Single-impurity XXZ circuit and prethermalisation

Having characterised transport via magnetisation-profile diagnostics in Secs. 3.3 and 3.5, we now switch to a complementary finite-time probe that ties directly to weakly broken conservation laws. We revisit the single-impurity Floquet XXZ circuit introduced in Secs. 2.7 and 2.7.2, where it served as a benchmark separating spectral chaos indicators from transport. Here we instead use it to isolate an *intensive* route to near-integrability: unlike the near-SWAP “swappy” regime of Sec. 3.4, where integrability is broken extensively by detuning all gates, the bulk is exactly integrable and only a *single* local defect violates the higher charges. This makes the prethermal mechanism especially transparent and allows us to identify a single slow rate from the decay of a representative conserved charge.

A standard route to prethermalisation in many-body systems is to start from an exactly solvable (typically integrable) reference dynamics and perturb it so that at least one conservation law is violated, but only weakly. In continuous time, this mechanism is often diagnosed by tracking the relaxation of local observables towards a *generalised* equilibrium of the reference model (a *GGE*), followed by a slow drift of the corresponding Lagrange multipliers towards a thermal Gibbs state. Two complementary viewpoints are widely used. (i) A *manifold-of-equilibria* picture: the state rapidly approaches an approximate GGE of the unperturbed model, and then evolves along this manifold under an effective autonomous equation with a rate controlled by the integrability-breaking perturbation [165]. (ii) A *kinetic* picture: one derives (or postulates) a Boltzmann-type evolution for quasi-particle occupations (or other slow mode variables) and identifies prethermal plateaux and their subsequent drift from the resulting scattering rates [161]. Both viewpoints motivate the same practical probes: (a) time dependence of putative conserved charges (or of their overlaps with observables), (b) correlators of those charges, and (c) extraction of a single slow rate (the smallest nonzero relaxation rate) governing the late part of the intermediate-time window.

Here we implement the same logic in discrete time, but with an *intensive* integrability-breaking perturbation: a single local defect inserted into an otherwise integrable U(1)-symmetric XXZ brickwork circuit. The resulting dynamics is particularly transparent because the perturbation is strictly local while the clean circuit carries an extensive family of local conserved charges.

3.7.1 Conserved charges of the clean XXZ circuit

For $h_{\text{imp}}=0$, the bulk brickwork circuit $\hat{U}_{\text{bulk}}$ is the integrable discrete-time XXZ Floquet chain introduced in [Sec. 2.5](#). Its local gate is the braid form of the unitary six-vertex R -matrix (written there as $\hat{u}(\lambda, \eta)$; [Eq. \(2.47\)](#)), and the associated commuting transfer matrices generate the extensive family of conserved charges mentioned in that section. To construct these charges explicitly, we follow [\[93\]](#) and embed the gate into a *unitary* six-vertex R -matrix $\check{R}(u)$ depending on a spectral parameter u and satisfying the braid Yang-Baxter equation

$$\check{R}_{12}(u-v) \check{R}_{23}(u) \check{R}_{12}(v) = \check{R}_{23}(v) \check{R}_{12}(u) \check{R}_{23}(u-v). \quad (3.28)$$

One then builds a staggered transfer matrix by tracing over an auxiliary spin- $\frac{1}{2}$ space,

$$\hat{T}(u) = \text{Tr}_a \left[\check{R}_{1a}(u+\tau) \check{R}_{2a}(u-\tau) \cdots \check{R}_{Na}(u+(-1)^{N-1}\tau) \right], \quad (3.29)$$

with staggering parameter τ (fixed by the choice of integrable bulk gate and not to be confused with the impurity strength h_{imp} introduced below). The Yang-Baxter relation implies the commuting family property

$$[\hat{T}(u), \hat{T}(v)] = 0 \quad \forall u, v, \quad (3.30)$$

and the homogeneous Floquet step $\hat{U}_{\text{bulk}}$ is obtained as $\hat{U}_{\text{bulk}} = \hat{T}(u_0)$ at the parameter value u_0 for which $\check{R}(u_0) = \hat{u}_{\text{XXZ}}(J, J_z)$ (up to an overall phase).

Conserved charges follow from the logarithmic-derivative expansion about $u = u_0$:

$$\hat{Q}_\ell^{(\pm)} = \left. \frac{\partial^\ell}{\partial u^\ell} \ln \hat{T}_\pm(u) \right|_{u=u_0}, \quad \ell = 1, 2, \dots, \quad (3.31)$$

where $\pm$ labels the two translation-inequivalent (even/odd-centred) staggerings, reflecting the odd/even substep structure of the brickwork update discussed in [Sec. 2.5](#). Each $\hat{Q}_\ell^{(\pm)}$ is extensive and has a strictly local density of finite range that increases with ℓ ; in particular $\ell = 1$ already has a three-site density [\[93\]](#). In what follows we fix

one choice of staggering and take

$$\hat{Q} = \hat{Q}_{\ell=1}, \quad (3.32)$$

orthogonalised against the trivial invariants $\hat{I}$ and total magnetisation $\hat{S}^z$, as the clean conserved quantity whose impurity-induced leakage we monitor.

3.7.2 Impurity Floquet step and autocorrelator probe




Figure 3.5: Prethermal decay of the charge autocorrelator $C_Q(t)$ in the single-impurity XXZ circuit. (a) Size dependence at $J = 1$ and $h_{\text{imp}} = 1$ for $J_z = 0.55, 1.1, 2.2$. (b) Anisotropy dependence at $N = 24$ and $h_{\text{imp}} = 1$. Evolved to $T = 10^6$ cycles.

We now introduce the *only* spatial inhomogeneity: a single on-site $\hat{z}$ -kick

$$\hat{u}_{\text{imp}} = \exp(-ih_{\text{imp}}\hat{s}_{n_0}^z), \quad n_0 = \left\lfloor \frac{N}{2} \right\rfloor, \quad (3.33)$$

applied once per Floquet period, yielding the single-impurity step $\hat{U}_{\text{SI}} = \hat{U}_{\text{bulk}} \hat{u}_{\text{imp}}$ (see Sec. 2.7.2 and Eq. (2.66)). This is the stroboscopic discretisation of the single-impurity XXZ setting studied in continuous time in Ref. [131], with the key difference that the bulk dynamics is a discrete-time integrable circuit rather than a Hamiltonian flow.

To probe the resulting near-integrable relaxation we use the infinite-temperature autocorrelator of $\hat{Q}$ in the $M=0$ symmetry sector,

$$C_Q(t) = \frac{\text{Tr}[\hat{Q}(t)\hat{Q}]}{\text{Tr}[\hat{Q}^2]}, \quad \hat{Q}(t) = \hat{U}_{\text{SI}}^{-t} \hat{Q} \hat{U}_{\text{SI}}^t, \quad t \in \mathbb{Z}. \quad (3.34)$$

For $h_{\text{imp}}=0$ one has $C_Q(t) = 1$. For $h_{\text{imp}} \neq 0$, $\hat{Q}$ is no longer conserved, but its violation is localised.

Fig. 3.5 shows three robust features. First, at fixed h_{imp} and varying anisotropy J_z/J ($J = 1$ and $J_z = 0.55, 1.1, 2.2$), $C_Q(t)$ displays an initial decay consistent with

hydrodynamic shedding of fast components (a non-universal, short-time regime), followed by a clean exponential window,

$$C_Q(t) \sim e^{-\Delta_N t}, \quad (3.35)$$

before eventually saturating to the finite- N floor set by the finite Hilbert space. The exponential regime is only weakly resolved at $J_z/J=0.55$ (suggesting limited scale separation at accessible sizes), whereas for $J_z/J = 1.1$ and 2.2 it is clean and extends across several decades in time. Second, because the impurity is *intensive* while $\hat{Q}$ is *extensive*, the leakage rate is expected to scale as $\Delta_N \propto N^{-1}$ (up to parameter-dependent prefactors): the perturbation only acts at $O(1)$ sites, while $\text{Tr}[\hat{Q}^2]$ grows extensively. This explains why plotting against the scaled time tJ/N collapses the pre-saturation curves across $N=16, \dots, 26$:

$$\Delta_N \approx \frac{\gamma(J, J_z, h_{\text{imp}})}{N}, \quad C_Q(t) \sim \exp\left[-\gamma \frac{t}{N}\right]. \quad (3.36)$$

The same finite-size collapse logic underlies resonance-based analyses of correlation decay in U(1) circuits.

Third, at fixed N (e.g. $N=24$) and varying h_{imp} , as shown in Fig. 3.5(b), the intermediate-time decay changes systematically, consistent with a defect-controlled leakage rate. Note that $h_{\text{imp}}=\pi$ is the maximum value of the impurity strength as Eq. (3.33) is 2π -periodic.

Finally, scanning the kick parameter J at fixed ratios (e.g. $J_z/J = 1.1$ and $h_{\text{imp}}/J = 1$) reveals a special role of the dual unitary (DU) point $J=\pi$, as shown in Fig. 3.6. At this point the bulk gate is DU and the stroboscopic response becomes exceptionally constrained; in the shown data this coincides with the onset of




Figure 3.6: Prethermal decay of $C_Q(t)$ in the single-impurity XXZ circuit at different J values, with $N=24$, $J_z/J = 1.1$, and $h_{\text{imp}}/J = 1$. The special role of the dual unitary (DU) point $J=\pi$ is evident: the stroboscopic response becomes exceptionally constrained and the exponential decay is replaced by a persistent subharmonic structure. Away from $J=\pi$ the same defect-controlled leakage mechanism governs the decay of $C_Q(t)$. At low J the exponential window is not well resolved, consistent with limited scale separation at accessible sizes. At large J still away from the dual unitary (DU) point, the exponential window is absent.

persistent subharmonic structure characteristic of a discrete-time-crystalline response. Away from $J=\pi$ the impurity-induced decay of $C_Q(t)$ resumes, governed by the same

defect-controlled leakage mechanism.

3.7.3 Leading resonance interpretation and the role of the Heisenberg time

A compact way to formalise the exponential window is to view the dynamics as the Heisenberg channel $\mathcal{U}_{\text{SI}}[\cdot] = \hat{U}_{\text{SI}}^\dagger(\cdot)\hat{U}_{\text{SI}}$ acting on operators and to restrict attention to the subspace of local (or quasi-local) operators. Projecting onto operators supported on at most r consecutive sites defines a truncated propagator $U_{\text{SI}}^{(r)} = \hat{P}_r \mathcal{U}_{\text{SI}} \hat{P}_r$ with spectrum inside the unit disk [166]. The eigenvalue of largest modulus below the fixed-point subspace,

$$|\lambda_*| = \max_{\lambda \notin (\text{fixed pts.})} |\lambda| < 1, \quad \Delta = -\ln |\lambda_*|, \quad (3.37)$$

acts as a leading resonance and yields the exponential regime (3.35) when $\hat{Q}$ has appreciable overlap with its eigen-operator. This is the discrete-time analogue of extracting the slowest relaxation rate from a kinetic equation [162] or from a flow on the GGE manifold [165].

The late-time plateau in finite- N numerics is not evidence of exact conservation. It is the generic finite-size saturation of autocorrelators in a finite Hilbert space. A clean observation of a prethermal decay therefore requires a separation between the prethermal lifetime $T \sim \Delta^{-1}$ and the Heisenberg time t_{H} of the relevant symmetry sector:

$$t_{\text{loc}} \ll t \ll T \ll t_{\text{H}}. \quad (3.38)$$

If T is too close to t_{H} at accessible sizes, the exponential window is truncated and the correlator appears to “stop decaying” early. This is the correct finite-size caveat when interpreting the saturation levels in Fig. 3.5(a) for small N .

While this resonance/Heisenberg-time framework provides a useful interpretation, we do not attempt a systematic extraction of λ_* from M_r nor a quantitative estimate of t_{H} here. A more detailed resonance analysis and its finite-size scaling therefore remains for further investigation.

3.8 Conclusion

Our results address transport across the phases of a highly generic $U(1)$ -symmetric disordered Floquet model, ranging from localised and ergodic regimes (sub-diffusive and diffusive, respectively) to a discrete-time “swappy” window in which excitations

propagate quasi-ballistically. To disentangle these behaviours we combined static and dynamic viewpoints within the $M=0$ sector: sector-resolved spectral probes computed with POLFED by applying the unitary gate-by-gate, and a local-quench protocol whose magnetisation profile is compressed into the first circular moment $R(t)$ (Sec. 3.2.4). The moment representation maps aggregated dynamics to simple trajectories in the complex plane and, because it is built solely from local Z -basis expectation values, is directly amenable to quantum-simulator experiments; its modulus and phase define spread and drift diagnostics, with the latter potentially connected to entanglement-speed notions near the DU line [167, 168].

Within this language the swappy regime is distinct and genuinely dynamical: it appears as an ergodic-swappy crossover near $J \simeq 2.1$ in the finite-time observables of Ch. 3 yet is absent from the static analyses of Sec. 2.1, and it is prethermal in the near-integrable sense, with crossover times that grow algebraically as $J \rightarrow \pi$ (roughly $t_* \propto (\pi - J)^{-2}$; Sec. 3.4 and Fig. 3.2). More broadly, these prethermal windows are not inherited from a continuous-time disordered XXZ Hamiltonian: the circuit is not a simple Trotterisation (the disorder does not scale with the kicking parameters), the protocol corresponds to strong static disorder periodically interrupted by finite unitary kicks, and the stroboscopic dynamics does not conserve energy. Consequently, prethermalisation is not organised by the conserved quantities of a parent Hamiltonian, but by quasi-conserved *Floquet* structures tied to solvable stroboscopic points and by the way they are broken. The single-impurity XXZ circuit makes this “digital matter” viewpoint concrete, an integrable bulk supports higher charges that are violated only locally, producing defect-controlled leakage and an exponential prethermal decay (Sec. 3.7 and Fig. 3.5), and, in the vicinity of $J=\pi$, a persistent subharmonic response consistent with a discrete-time-crystalline behaviour.

3.9 Methods

3.9.1 Efficient time evolution of a state by a random circuit

The numerical method we developed circumvents the need to construct the entire $\hat{U}$ matrix. Instead, we apply each gate directly to the state vector, leveraging the fact that both static and dynamic analyses necessitate the application of the Floquet operator to a state. This allows the memory cost of diagonalising $\hat{U}$ to be $\mathcal{O}(d)$ and the number of operations to be $\mathcal{O}(Nd)$, similar to what pointed out in Ref. [115].

Our approach involves recording the action of a gate on the first two spins of a state $|\psi\rangle$. If a gate acts on spins a and $a+1$, we shift the system by $-a$ sites, apply the gate to the first two spins, and then shift the system back to its original configuration.

To implement this, we require two key components. First, we must determine the value of the first two spins for each element of $|\psi\rangle$. The state is represented as a d -vector:

$$|\psi\rangle = \sum_i c_i |i\rangle \quad (3.39)$$

where the i -th element corresponds to the $|i\rangle$ state. Hence, we need to identify which segments of the d -vector have the first and second spins up, namely which of the c_i are associated with a $|i\rangle$ with first and second spin up.

The second ingredient involves the transformation $f_{(1)}$, which maps a state from sites $0, \dots, N-1$ to the corresponding state on sites $N-1, 0, \dots, N-2$. By recursively applying $f_{(1)}$, we can achieve the required shift, i.e., $f_{(a)} = f_{(1)}^a$. This strategy offers a significant memory advantage, as it only necessitates storing vectors of size d . Although this comes with a cost in the number of operations, each gate of $\hat{U}$ or $\hat{K}$ must be applied sequentially, the memory efficiency of this method allows for better parallel computation of multiple trajectories, which is crucial for disorder averaging.

3.9.2 Numerical methods: vector propagation and typicality in the $M=0$ sector

All diagnostics in [Secs. 3.3](#) and [3.4](#) are computed by propagating state vectors in the symmetry-reduced Hilbert space $\mathcal{H}^{(0)}$ and measuring local observables. No density matrix is stored.

One realisation and stroboscopic sampling

A single realisation consists of: (i) sampling disorder phases $\{h_n, h'_n, \phi_n\}$ for each two-site gate, (ii) sampling a bond permutation P defining the circuit order in [Eq. \(3.19\)](#), and (iii) preparing one projected typical state implementing [Eq. \(3.4\)](#) (see below). Time evolution is generated by repeated application of $\hat{U}$ to a state vector,

$$|\psi(t)\rangle = \hat{U}^t |\psi(0)\rangle. \quad (3.40)$$

We record the dynamics cycle-by-cycle for early times (in practice, the first $t=100$ cycles). At later times we may optionally sample more coarsely by applying blocks of Floquet cycles between measurements. This reduces storage/IO costs without changing the underlying dynamics; all fits and summaries are taken on time windows where the sampling is sufficiently fine to resolve the drift and spreading diagnostics.

Typicality and the projected typical initial state

Let $d_0 = \dim \mathcal{H}^{(0)} = \binom{N}{N/2}$. For a typical random state $|\psi_{\text{rand}}^{(0)}\rangle \in \mathcal{H}^{(0)}$, expectation values approximate the infinite-temperature trace in the same sector:

$$\langle \psi_{\text{rand}}^{(0)} | \hat{A} | \psi_{\text{rand}}^{(0)} \rangle \simeq \frac{1}{d_0} \text{Tr}[\hat{P}_0 \hat{A}] \quad (3.41)$$

with fluctuations suppressed as $\mathcal{O}(d_0^{-1/2})$. We implement the mixed initial condition [Eq. \(3.4\)](#) by projecting $|\psi_{\text{rand}}^{(0)}\rangle$ with $\hat{P}_{n_0}^\dagger$:

$$|\psi(0)\rangle = \frac{\hat{P}_{n_0}^\dagger |\psi_{\text{rand}}^{(0)}\rangle}{\sqrt{\langle \psi_{\text{rand}}^{(0)} | \hat{P}_{n_0}^\dagger | \psi_{\text{rand}}^{(0)} \rangle}}. \quad (3.42)$$

The magnetisation profile is then measured as

$$M_n(t) = \langle \psi(t) | \hat{s}_n^z | \psi(t) \rangle. \quad (3.43)$$

From $M_n(t)$ we first form $p_n(t)$ via [Eq. \(3.6\)](#) and then construct $R(t)$ via [Eq. \(3.9\)](#). The diagnostics $\mu(t)$, $\sigma(t)$, and $\nu(t)$ follow from [Eqs. \(3.12\)](#) and [\(3.15\)](#).

The explicit derivation connecting the profile protocol [Eqs. \(3.42\)](#) and [\(3.43\)](#) to sector-restricted dynamical correlators (and to stochastic trace estimation) is technical and is therefore deferred to [Sec. C](#).



Chapter 4

Anomalous equilibration: Mpemba effect

The Mpemba effect is a counterintuitive relaxation phenomenon: under the same dynamics, the further a system is initially prepared from equilibrium, the faster it reaches it. In its original form, this corresponds to the anomalous situation of hot water cooling (or freezing) faster than warm water, first investigated by Mpemba and Osborne in 1969 [169]; the same phenomenology is now known to be ubiquitous across nonequilibrium relaxation processes [170]. It has since been observed in various physical systems ranging from clathrate hydrates [171] to crystallisation processes in polymers [172] and magnetic transitions in alloys [173]. Following its precise mathematical characterisation in classical systems weakly coupled to Markovian environments [174], the effect was further investigated theoretically [170, 175–177] and experimentally [178–180] for colloidal systems. Two further important directions in the study of the classical Mpemba effect concern driven granular gases [170, 181, 182] and phase transitions [183, 184]. The latter, despite pertaining to the phenomenon as originally observed in water, remains the most challenging and least understood case [185, 186].

In recent years, the theoretical framework was extended to quantum systems [187–198]. In this context, a related phenomenon has been observed in the dynamical restoration of symmetries: states that initially display stronger symmetry breaking can locally¹ restore it faster than states with initially weaker symmetry breaking [199, 200]. We refer to this phenomenon as the symmetry Mpemba effect, to differentiate it from the conventional thermal Mpemba effect. A notable recent development is the Pontus-Mpemba effect [201], which shows that even for a fixed initial state a two-step protocol can outperform direct relaxation: a deliberate random “heating” stage that drives the system further from equilibrium, followed by free relaxation, can accelerate the approach to equilibrium compared to evolving freely from the start.

In this chapter, we establish a framework to understand both the thermal and symmetry Mpemba effects on the same footing using resource theories (RTs) [202, 203]. A Mpemba effect arises when two resourceful states evolve under the same dynamics, and the more

¹Throughout this chapter we consider dynamics on a composite Hilbert space $\mathcal{H} = \mathcal{H}_s \otimes \mathcal{H}_e$, where $\mathcal{H}_s$ (system) and $\mathcal{H}_e$ (environment) are always distinguished. In the open-dynamics case $\mathcal{H}_e$ describes an external bath; in the closed-dynamics case $\mathcal{H}_e$ is the complementary subsystem within an overall isolated $\mathcal{H}$. Local symmetry restoration then means that, after evolving under the joint dynamics, we apply the partial trace Tr_e and study symmetry recovery on the reduced state $\hat{\rho}_s = \text{Tr}_e[\hat{\sigma}]$.




Figure 4.1: In a resource-theoretic framework, the Mpemba effect occurs when a state that initially possesses more of a given resource depletes that resource faster than a less resourceful state, under the evolution by the same free operation, so that their resource monotones cross. This single picture unifies a variety of anomalous equilibration phenomena (for example, restoring thermal equilibrium or symmetry in classical and quantum systems). Mpemba physics then becomes the study of why different initial states dissipate resources at different rates, and how we can harness those differences to engineer exotic effects such as ultrafast cooling. In this chapter, we apply this analysis to the specific resource theories shown in the schematic.

resourceful one dissipates the resource faster than the less resourceful one, causing their resource monotones to cross (Fig. 4.1). This description allows us to capture the essence of both effects within a single conceptual framework, and it naturally extends to other resource notions. Building upon this, we employ the framework of modes of asymmetry [204] to explain the occurrence of the symmetry Mpemba effect. Analogously to how the thermal Mpemba effect is understood in terms of overlaps with slow relaxation modes [174], we demonstrate that the symmetry Mpemba effect also arises when strongly symmetry-broken initial states have small (or vanishing) overlap with the slowest symmetry-restoring mode. Within this framework, we study various examples revealing that the thermal Mpemba effect can also occur in subsystems of unitarily evolving systems, while the symmetry Mpemba effect can manifest in systems governed by Lindblad dynamics, as well as circuits with abelian (continuous) and non-abelian symmetries. In particular we find the first example of symmetry Mpemba effect in a classical system. In this context, we emphasise that the specific type of Mpemba effect observed is dictated by the chosen measure for quantifying the relaxation process, rather than being an intrinsic property of the physical system itself. Finally, in the presence of arbitrary symmetries, we highlight the connection between the non-equilibrium free energy, commonly used for the thermal Mpemba effect, and the relative entropy of asymmetry (also known as entanglement asymmetry), a measure for the asymmetry of a state.

We then push the resource-theoretic picture beyond thermodynamic and symmetry settings by studying Mpemba physics for *complexity resources*, imaginarity, non-Gaussianity, and magic, in closed many-body dynamics generated by resource-free random circuits. In this setting, resources are conserved *globally* under free unitaries yet can dissipate *locally* on a subsystem, giving rise to quantum Mpemba and Pontus-Mpemba effects for abstract measures of quantum complexity. To guide the reader

through the chapter, Mpemba effects originating in different RTs are indicated via different symbols: $\square$ for athermality, $\circ$ for asymmetry, and $\triangle$ for symmetry-respecting athermality, $\heartsuit$ for imaginarity, $\diamond$ for (fermionic) non-Gaussianity, and $\star$ for magic.

This chapter is organised as follows. We begin with a technical aside collecting the explicit Davies generator used in the open-system examples (Sec. 4.1). In Sec. 4.2 we review the basic resource-theoretic notions (free states/operations and monotones) that underpin our approach. We then study thermal Mpemba effects in classical and quantum relaxation (Sec. 4.3) and symmetry Mpemba effects via modes of asymmetry (Sec. 4.4). In Sec. 4.5 we connect these viewpoints and relate free-energy and asymmetry-based diagnostics within a single relative-entropy decomposition. Finally, Sec. 4.6 extends the analysis to complexity resources in unitary circuits and introduces Pontus-Mpemba protocols as a systematic speedup mechanism, before we summarise and discuss outlooks in Sec. 4.7.

4.1 Details on the Davies map

The Davies map arises from a particular Lindblad equation describing the thermalisation of quantum systems weakly coupled to a Markovian heat bath [205]. In the main text, we have focused on its general mathematical properties, while here we will give the explicit expression of the equation that we used to compute various examples. Consider the spectral decomposition of the system's Hamiltonian $\hat{H}_s = \sum_k h_k |k\rangle\langle k|$. In the eigenbasis of $\hat{H}_s$, the Davies generator reads

$$\begin{aligned} \frac{d\hat{\rho}_s(t)}{dt} = \mathcal{L}[\hat{\rho}_s(t)] = & -i[\hat{H}_s, \hat{\rho}_s(t)] \\ & + \sum_{k < k'} \mathcal{D}_{kk'}^{(1)}[\hat{\rho}_s(t)] + \mathcal{D}_{kk'}^{(2)}[\hat{\rho}_s(t)], \end{aligned} \quad (4.1)$$

where the dissipators $\mathcal{D}_{kk'}^{(1)}$ and $\mathcal{D}_{kk'}^{(2)}$ are generated by the jump operators $\hat{L}_{kk'}^{(1)}$ and $\hat{L}_{kk'}^{(2)}$, respectively, which couple all Hamiltonian eigenstates as

$$\begin{aligned} \hat{L}_{kk'}^{(1)}(\beta) &= (1 \mp f^\pm(\beta, h_{k'} - h_k))^{1/2} |k\rangle\langle k'|, \\ \hat{L}_{kk'}^{(2)}(\beta) &= (f^\pm(\beta, h_{k'} - h_k))^{1/2} |k'\rangle\langle k|. \end{aligned} \quad (4.2)$$

Here $f^\pm(\beta, h_{k'} - h_k) = 1/(\exp(\beta(h_{k'} - h_k)) \pm 1)$ are the Fermi (+) and the Bose (−) distribution functions which enforce the detailed balance condition, ensuring that the unique steady state is the thermal state at inverse temperature β . Throughout all examples in the main text, we considered bosonic canonical commutation relations.

4.2 Resource Theories

Resource theories (RTs) (see Ref. [203] for a review) provide the natural language for understanding and unifying diverse Mpemba effects: they formalise which state transformations are possible when only certain free operations are allowed, and thereby translate relaxation processes into the decay of a resource monotone [206]. In a nutshell, a resource theory is defined by two foundational ingredients: a set of free states and a class of free operations. Free operations are completely positive and trace-preserving (CPTP) maps that are considered accessible at no cost, either because they are physically implementable with minimal resources or because they reflect fundamental constraints of the theory. Similarly, the set of free states, denoted $\mathcal{F} \subset \mathcal{S}(\mathcal{H}_s)$ (where $\mathcal{S}(\mathcal{H}_s)$ indicates the set of density operators on the Hilbert space $\mathcal{H}_s$), consists of states that can be prepared at no cost. A minimal requirement that any resource theory must satisfy is that the set of free states is closed under free operations (as defined later in this section); that is, for any free state $\hat{\pi} \in \mathcal{F}$ and any free operation $\mathcal{E}$, we have $\mathcal{E}[\hat{\pi}] \in \mathcal{F}$. This closure condition ensures that no resource can be generated from within the free set. All states outside the free set, i.e., $\mathcal{S}(\mathcal{H}_s) \setminus \mathcal{F}$ ², are called resourceful or non-free. These are states that cannot be obtained from any free state using only free operations. Their preparation requires access to non-free operations, and they are typically the carriers of some operational advantage within the theory. The presence and quantification of such states lie at the core of resource-theoretic approaches.

To introduce these concepts, we recall two prominent examples of resource theories: bipartite entanglement and athermality in classical systems. In the theory of bipartite entanglement, the free states are separable states, and the free operations are local operations and classical communication (LOCC) [207]. Entangled states lie outside the free set and enable tasks such as quantum teleportation or the violation of Bell inequalities, which are impossible under LOCC alone. A fully classical example is provided by the RT of athermality for classical systems. Here, one considers a system in contact with a thermal environment at inverse temperature β . The free operations, also called thermal operations (TOs), are stochastic maps with fixed point given by the Gibbs distribution $(\pi_\beta)_i \propto e^{-\beta E_i}$, which is also the only free state [208]. Any state that departs from π_β therefore possesses athermality, representing a source of non-equilibrium free energy that can be harnessed to perform useful work or drive transformations that are otherwise forbidden under thermal operations. We analyse this theory in more detail in Sec. 4.3.1, showing how it naturally extends to quantum systems too.

A resource monotone is a function $M : \mathcal{S}(\mathcal{H}_s) \rightarrow [0, \infty)$ satisfying two key properties:

²We indicate by $\mathcal{S}(\mathcal{H}_s)$ the set of density matrices on $\mathcal{H}_s$.

1. Vanishing for free states: $\hat{\rho} \in \mathcal{F} \Rightarrow M(\hat{\rho}) = 0$.
2. Monotonicity under free operations: for every state $\hat{\rho}$ and every $\mathcal{E} \in \mathcal{O}$,

$$M(\mathcal{E}[\hat{\rho}]) \leq M(\hat{\rho}), \quad (4.3)$$

where the set of free operations $\mathcal{O}$ consist of all CPTP maps $\mathcal{E} : \mathcal{H}_s \rightarrow \mathcal{H}_s$ that preserve the free set.

A Mpemba effect manifests when two initial states $\hat{\rho}_1$ and $\hat{\rho}_2$ satisfying $M(\hat{\rho}_1) > M(\hat{\rho}_2)$ evolve under the same free operation $\mathcal{E}_t$ so that

$$M(\mathcal{E}_t[\hat{\rho}_1]) < M(\mathcal{E}_t[\hat{\rho}_2]) \quad \text{for some } t > 0, \quad (4.4)$$

indicating that the more resourceful state has dissipated the resource faster.

To quantify resourcefulness, we adopt the family of Petz-Rényi α -divergences [209]

$$M_\alpha(\hat{\rho}) = \min_{\hat{\pi} \in \mathcal{F}} S_\alpha(\hat{\rho} \parallel \hat{\pi}), \quad \alpha \in (0, 1) \cup (1, 2], \quad (4.5)$$

where

$$S_\alpha(\hat{\rho} \parallel \hat{\pi}) = \frac{1}{\alpha - 1} \ln \text{Tr}[\hat{\rho}^\alpha \hat{\pi}^{1-\alpha}], \quad (4.6)$$

obeys the data-processing inequality under any CPTP map [203, 209] and recovers the quantum relative entropy of resource for $\alpha \rightarrow 1$, i.e.,

$$M(\hat{\rho}) = \min_{\hat{\pi} \in \mathcal{F}} S(\hat{\rho} \parallel \hat{\pi}), \quad (4.7)$$

where $S(\hat{\rho} \parallel \hat{\pi}) = \text{Tr}[\hat{\rho}(\ln \hat{\rho} - \ln \hat{\pi})]$ is the quantum relative entropy, i.e., the quantum generalisation of the Kullback-Leibler (KL)-divergence. Physically, M_α measures the divergence of $\hat{\rho}$ with respect to the closest free state in the Rényi sense. We note that Eqs. (4.5) and (4.7) follow the standard entropic/divergence route to construct monotones in general resource theories (vanish on the free states and nonincreasing under the free operations, see Ref. [210]).

4.3 The thermal Mpemba effect

The first rigorous mathematical description of the thermal Mpemba effect was given in 2017 by Lu and Raz, who studied discrete-state continuous-time Markovian processes. In [174], the dynamics is generated by a classical Liouvillian $\hat{\mathcal{L}}^{\text{cl}}$, whose eigenvalues λ_k

are non-positive, and it is convenient to order them in decreasing order of their real part as $0 = \lambda_1 > \text{Re}(\lambda_2) \geq \text{Re}(\lambda_3) \geq \dots$. The classical Liouville equation (see [Sec. 4.3.2](#)) dictates that the contribution to the dynamics from all eigenmodes $\mathbf{r}_k$ of $\hat{\mathcal{L}}^{\text{cl}}$ decay exponentially in time with a rate given by the real part of the corresponding eigenvalue, except for $\mathbf{r}_1$ corresponding to $\lambda_1 = 0$ that describes the thermal steady state. Thus, at long times, the equilibration speed of typical states will be proportional to $\exp(\text{Re}(\lambda_2)t)$. Instead, special initial states having zero overlap with the so-called slowest-decaying mode $\mathbf{r}_2$ will equilibrate exponentially faster, with speed $\propto \exp(\text{Re}(\lambda_3)t)$. Moreover, if a faster equilibrating state happens to be initially further away from the steady state than a typical state, then the distance of the two states measured with some function M will cross in time and this is called a strong Mpemba effect. If the fast equilibrating states have small but non-vanishing overlap with the slowest-decaying mode, the speedup will not be exponential, and the crossing indicates a weaker Mpemba effect. To set the stage for the following examples, we first review athermality, the RT relevant to the thermal Mpemba effect.

4.3.1 Athermality as a Resource

Athermality, as anticipated for classical system in [Sec. 4.2](#), a state's deviation from thermal equilibrium $\hat{\pi}_\beta$, carries useful non-equilibrium free energy $\Delta F(\hat{\rho}) = F(\hat{\rho}) - F(\hat{\pi}_\beta)$, where $F(\hat{\rho}) = \text{Tr}[\hat{H}_s \hat{\rho}] - \beta^{-1} S(\hat{\rho}) = \beta^{-1} S(\hat{\rho} \| \hat{\pi}_\beta)$ [211]. In the resource theory of athermality, it quantifies the resource that can be converted into work using thermal operations [212, 213] (defined more precisely below). These consist of operations that can be realised by coupling the system to an ancilla prepared in the Gibbs state at inverse temperature β , applying an energy-preserving unitary, and tracing out the ancilla afterwards, allowing heat exchange while preserving overall energy conservation. Framing the thermal Mpemba effect in this RT underscores what drives the equilibration process for the thermal Mpemba effect, even though, as with any non-equilibrium process, this free energy ultimately dissipates irreversibly into the bath rather than being fully harvested as work [214, 215].

In this framework, the only free state is the thermal state at the inverse temperature β set by the environment (throughout this chapter, we will consider $k_B = \hbar = 1$); any other state is considered resourceful. A quantum channel $\mathcal{E} : \mathcal{B}(\mathcal{H}_s) \rightarrow \mathcal{B}(\mathcal{H}_s)$ (where $\mathcal{B}(\mathcal{H}_s)$ denotes the set of bounded operators on $\mathcal{H}_s$) is a free operation, denoted TO, if there exists [216]

- a thermal bath in a Gibbs state $\hat{\pi}_\beta = e^{-\beta \hat{H}_e} / Z_e$ with arbitrary Hamiltonian $\hat{H}_e$ and fixed inverse temperature β ,
- a global energy-conserving unitary $\hat{U} \in \mathcal{B}(\mathcal{H}_s \otimes \mathcal{H}_e)$ such that $[\hat{U}, \hat{H}_s \otimes \hat{H}_e] = 0$,

where the action of $\mathcal{E}$ on a system state $\hat{\rho}$ is given by

$$\mathcal{E}[\hat{\rho}] = \text{Tr}_e[\hat{U}(\hat{\rho} \otimes \hat{\pi}_\beta)\hat{U}^\dagger]. \quad (4.8)$$

Trivially, Eq. (4.5) becomes

$$M_\alpha(\hat{\rho}(t)) = S_\alpha(\hat{\rho}_s(t) \parallel \hat{\pi}_\beta), \quad (4.9)$$

which for $\alpha \rightarrow 1$ recovers the relative entropy of athermality

$$M(\hat{\rho}(t)) = S(\hat{\rho}_s(t) \parallel \hat{\pi}_\beta). \quad (4.10)$$

For states that are diagonal in the energy eigenbasis, (i.e., classical states), this reduces to the KL-divergence [217]. An important property is that a state $\hat{\rho}$ can be transformed into another state $\hat{\sigma}$ via TOs only if $\hat{\rho}$ thermomajorises $\hat{\sigma}$ [218]. Notably, it was shown that for classical Markovian thermal operations, thermomajorisation can be used to determine whether the Mpemba effect occurs for all monotone functions [219]. In the next section, we will explore examples of classical and quantum thermal Mpemba effects from the perspective of the resource theory of athermality. More generally, one can consider open system dynamics generated by a Lindbladian $\mathcal{L}$ whose unique steady state $\hat{\pi}$ need not be Gibbsian, as often studied in the context of the Mpemba effect [188, 190, 192, 220]. In that setting, the relevant resource theory is non-stationarity, which we will return to in Sec. 4.5. To illustrate the role of athermality in the thermal Mpemba effect, we now turn to two paradigmatic examples: a classical system akin to that in Ref. [174], and a quantum counterpart inspired by Ref. [194]. In Sec. 4.3.4, we instead demonstrate how this phenomenon can arise in isolated settings. In Ref. [221], an alternative resource-theoretic interpretation of the thermal Mpemba effect is developed, using thermomajorisation to analyse the roles of correlations, coherence, non-Markovianity, and dimensionality in both classical and quantum relaxation dynamics, with water serving as a central example.

4.3.2 Example 1: Thermalisation of a Spin Chain

The quantum resource-theoretic framework for athermality introduced above can also be applied to classical dynamics, such as for instance discrete-state, continuous-time Markov processes [170, 174, 222]. In this context, the monotone for athermality Eq. (4.9), i.e. the KL-divergence, has already been studied. The KL-divergence serves as a generalised free energy: it is non-increasing under any stochastic map preserving the steady state π_β , and its decay is due to the depletion of classical athermality.

Consider a classical system whose state is specified by a probability distribution

$$\mathbf{p}(t) = (p_1(t), p_2(t), \dots, p_L(t))^\top, \quad (4.11)$$

where the subscript $i = 1, 2, \dots, L$ labels discrete microstates. For the classical Markovian Liouvillian dynamics

$$\frac{d\mathbf{p}(t)}{dt} = \hat{\mathcal{L}}^{\text{cl}} \mathbf{p}(t), \quad (4.12)$$

the fixed point is the thermal distribution $\mathbf{p}(t \rightarrow \infty) = \boldsymbol{\pi}_\beta$ at the inverse temperature β if the classical Liouvillian matrix $\hat{\mathcal{L}}^{\text{cl}}$ is

$$\hat{\mathcal{L}}_{ii'}^{\text{cl}} = \begin{cases} W_{ii'}, & i \neq i', \\ -\sum_{i \neq i'} W_{ii'}, & i = i'. \end{cases} \quad (4.13)$$

Here, $W_{ii'}$ satisfies the classical detailed balance condition $W_{ii'}/W_{i'i} = e^{-\beta(E_i - E_{i'})}$, where E_i denotes the energy of the i -th microstate. Note that the diagonal terms ensure probability conservation in the system. It can be shown that $\hat{\mathcal{L}}^{\text{cl}}$ is Hermitian with respect to the $\boldsymbol{\pi}_\beta$ -weighted inner product, and it is useful to sort its eigenvalues λ_k and its corresponding eigenmodes $\mathbf{r}_k$ as $0 = \lambda_1 > \lambda_2 > \dots$. Following Ref. [174], we study a classical spin-1/2 chain, where an upwards and a downwards pointing spin on site n is described by $s_n = +1$ and $s_n = -1$, respectively. We study N_s spins interacting locally with their nearest neighbours with coupling J and with a homogeneous external magnetic field with interaction strength h , for which the energy of a microstate reads

$$E_i = -J \sum_{n=1}^{N_s} s_n s_{n+1} - h \sum_{n=1}^{N_s} s_n, \quad (4.14)$$

where the microstate index $i \in \{0, 1, \dots, 2^{N_s} - 1\}$ uniquely labels the configuration $(s_1, s_2, \dots, s_{N_s})$ via binary encoding. We consider fixed boundary conditions $s_{N_s+1} = s_1 = +1$ and study the Liouvillian dynamics in Eq. (4.12) by plugging Eq. (4.14) into Eq. (4.13) under the constraint of allowing only transitions between microstates i and i' that differ by a single spin flip. Different from [174], for a given initial probability distribution $\mathbf{p}(t_0)$, we construct the fastest-thermalising, unitarily-connected state. This is the permutation of the initial $\mathbf{p}$ that minimises the overlap with the slowest-decaying modes of $\hat{\mathcal{L}}^{\text{cl}}$. Since the number of permutations grows factorially with the number of microstates, we resort to a Metropolis algorithm outlined in [194] to minimise the cost function

$$C = \sum_{k=2}^{K+1} |\langle \mathbf{p}, \mathbf{r}_k \rangle_{\boldsymbol{\pi}_\beta}|, \quad (4.15)$$

describing the overlap of the initial state with $K < L - 1$ slowest-decaying modes, with

$$\langle \mathbf{p}, \mathbf{r}_k \rangle_{\boldsymbol{\pi}_\beta} = \sum_i \frac{p_i(\mathbf{r}_k)_i}{(\boldsymbol{\pi}_\beta)_i} \quad (4.16)$$

being the so- called Markovian inner product [223].

In short, at every iteration, we perform stochastic permutations of four randomly selected spins and always accept the new configuration if the cost function C' is decreased. To avoid local minima, if $C' > C$, we still accept the new configuration with a probability $p_{\text{acc}} = \exp(-T_{\text{eff}}(C' - C))$, where the effective inverse temperature T_{eff} is decreased linearly with the number of successful iterations. Note that such optimised (non-equilibrium) initialisations are significantly more challenging to prepare experimentally, since, unlike thermal states, they require control over individual spins [224].

For the thermal operation generated by $\hat{\mathcal{L}}^{\text{cl}}$, the monotone defined in Eq. (4.10) simplifies to the KL-divergence

$$M(\mathbf{p}(t)) = D(\mathbf{p}(t) \parallel \boldsymbol{\pi}_\beta) = \sum_i p_i(t) \ln \left(\frac{p_i(t)}{(\boldsymbol{\pi}_\beta)_i} \right), \quad (4.17)$$

which measures the divergence of $\mathbf{p}(t)$ with respect to the thermal state. The state can be decomposed as

$$\mathbf{p}(t) = \sum_{k \geq 1} a_k \mathbf{r}_k e^{\lambda_k t}, \quad a_k = \langle \mathbf{p}, \mathbf{r}_k \rangle_{\boldsymbol{\pi}_\beta}, \quad (4.18)$$

and asymptotically, the behaviour of the KL-divergence will be

$$D(\mathbf{p}(t) \parallel \boldsymbol{\pi}_\beta) \sim \frac{1}{2} |a_2|^2 e^{2\text{Re}(\lambda_2)t}, \quad (4.19)$$

as shown in Sec. F. In other words, the KL-divergence decays exponentially with a rate being twice the spectral gap $-\text{Re} \lambda_2$. When the initially more resourceful state has zero overlap with the slowest decaying mode, we are guaranteed to find a Mpemba crossing. Instead, Eq. (4.19) allows us to assess whether a crossing will also happen in the context of a weak Mpemba effect (i.e. diminished but not vanishing overlap with the slowest decaying mode for the initially more resourceful state).




Figure 4.2: Thermal Mpemba effect in an open classical spin chain $\square$. We plot the KL divergence $D(\mathbf{p}(t) \parallel \boldsymbol{\pi}_\beta)$ for Eq. (4.14) during relaxation from $\beta' = 0.5$ to a bath at $\beta = 1$ ($J = -0.4$, $h = 0.2$, $N_s = 7$). Insets: crossing time τ_α of D_α Eq. (4.20) vs α (top right) and Liouvillian eigenvalues near 0, with λ_2 (thermal) and λ_4 (optimised) setting the asymptotic decay (bottom left). Optimised state: $C \simeq 2 \times 10^{-5}$ in Eq. (4.15).

In Fig. 4.2, we show the dynamics of the KL-divergence during a cooling process for a

thermal state at inverse temperature β' (full line) and an optimised state (dashed line) whose overlap with the two slowest-decaying modes of the Liouvillian was minimised. The optimised state is initially further away from the equilibrium state but thermalises exponentially faster, displaying an Mpemba crossing $\square$.

Similarly to Eq. (4.9), the KL-divergence can be generalised to a one-parameter family of functions that are also non-increasing under Markovian dynamics, known as the Rényi relative entropies [225]

$$M_\alpha(\mathbf{p}(t)) = D_\alpha(\mathbf{p}(t) \parallel \boldsymbol{\pi}_\beta) = \frac{1}{\alpha - 1} \ln \left(\sum_i p_i^\alpha (\boldsymbol{\pi}_\beta)_i^{1-\alpha} \right), \quad (4.20)$$

which correspond to the classical version of Eq. (4.9). The upper-right inset of Fig. 4.2 shows that the Mpemba crossing time τ_α between the thermal and optimised states increases monotonically with α . This illustrates that τ_α lacks intrinsic physical significance, as it depends on the choice of divergence measure. In particular, the variation of τ_α with α implies that different α -divergences can yield qualitatively different conclusions: for certain values of α , no crossing may be detected. We stress that this argument relies on allowing for general non-equilibrium initial states.

4.3.3 Example 2: The Davies Map

The generalisation of Eq. (4.12) to quantum systems is given by the Lindblad equation

$$\frac{d\hat{\rho}_s(t)}{dt} = \mathcal{L}[\hat{\rho}_s(t)]. \quad (4.21)$$

Here the superoperator $\mathcal{L} = \mathcal{H}_s + \mathcal{D}$ evolving the state $\hat{\rho}_s$ consists of a unitary part $\mathcal{H}_s[\hat{\rho}_s(t)] = -i[\hat{H}_s, \hat{\rho}_s(t)]$, where $\hat{H}_s$ denotes the system Hamiltonian, and a dissipative part

$$\mathcal{D}[\hat{\rho}_s(t)] = \sum_l \hat{L}_l \hat{\rho}_s(t) \hat{L}_l^\dagger - \frac{1}{2} \left\{ \hat{L}_l^\dagger \hat{L}_l, \hat{\rho}_s(t) \right\}, \quad (4.22)$$

with the jump operators $\hat{L}_l$ describing the effect of an environment on the system. Here we focus on a Lindbladian describing the weak interaction of a system with a Markovian heat bath, known as the Davies map [205], which is an important example of an elementary Markovian quantum thermal operation [226]. The Davies generator is characterised by two mathematical properties: it satisfies the quantum detailed balance condition [227] and the superoperators corresponding to the unitary and the dissipative part of its generator commute. These properties imply that the unique fixed point of the Davies map is the thermal state $\hat{\pi}_\beta = e^{-\beta \hat{H}_s} / \text{Tr}[e^{-\beta \hat{H}_s}]$. Moreover, in the energy eigenbasis, the Davies generator can be written in a block-diagonal form, where one block describes the thermalisation of the populations and the other the

thermalisation of the coherences, which evolve independently from one another (as we will analyse thoroughly in [Sec. 4.4.3](#)). In [\[194\]](#), the block structure of the Davies map




Figure 4.3: Thermal Mpemba effect in an open quantum system. Single-qubit thermalisation via the Davies map [\[205\]](#): the unitarily optimised/rotated state of [Ref. \[194\]](#) (dashed) shows faster exponential decay of $M(\hat{\rho}(t))$ ([Eq. \(4.24\)](#)) than a random initial state (solid), yielding an Mpemba crossing. Insets: crossing time τ_α of M_α ([Eq. \(4.5\)](#)) vs α (top right) and Davies-generator spectrum including complex eigenvalues (bottom left). Parameters: $h = 10$, $\beta = 0.1$, $\mathbf{r} = (0.221, 0.867, 0.206)$, $\mathbf{r}' = (0, 0, 0.919)$.

was exploited to investigate the thermal Mpemba effect. In particular, it was shown that for any initial state with non-vanishing coherences in the energy eigenbasis, one can always construct an exact unitary transformation $\hat{R}_{\text{coh}}$ that leads to exponentially faster thermalisation, provided that the Liouvillian gap is defined by an eigenvalue with non-vanishing imaginary part, i.e. $\Im(\lambda_2) \neq 0$. At the same time, a population inversion of the rotated state, implemented by another unitary transformation $\hat{R}_{\text{inv}}$, always increases the state's non-equilibrium free energy

$$F(\hat{\rho}) = \text{Tr}[\hat{\rho} \hat{H}] - \beta^{-1} S(\hat{\rho}). \quad (4.23)$$

A key aspect of this fundamental thermodynamic quantity is that its (positive) deviation from the equilibrium free energy is proportional to the relative entropy of athermality [\[211\]](#)

$$\Delta F(t) = \beta^{-1} M(\hat{\rho}_s(t)) = \beta^{-1} (F(\hat{\rho}_s(t)) - F(\hat{\pi}_\beta)) = \beta^{-1} S(\hat{\rho}_s(t) \parallel \hat{\pi}_\beta), \quad (4.24)$$

where the M is the monotone defined in [Eq. \(4.10\)](#). Crucially, [Eq. \(4.24\)](#) decreases monotonically during the Davies dynamics (as seen in [Sec. 4.1](#)).

By decomposing the state in the eigenmodes of $\mathcal{L}$ similarly to [Eq. \(4.18\)](#), one finds

$$\hat{\rho}(t) = \sum_{k \geq 1} a_k \hat{r}_k e^{\lambda_k t}, \quad a_k = \text{Tr}[\hat{\ell}_k^\dagger \hat{\rho}], \quad (4.25)$$

which for long times (as we derive in [Sec. F](#)) implies

$$S(\hat{\rho}(t) \parallel \hat{\pi}_\beta) \sim \frac{1}{2} |a_j|^2 e^{2\text{Re}(\lambda_j)t} \text{Tr}[\hat{r}_j^\dagger \hat{\pi}_\beta^{-1} \hat{r}_j], \quad (4.26)$$

where $j = \arg \max_{j: a_j \neq 0} \text{Re} \lambda_j$ labels the slowest-decaying mode of $\mathcal{L}$ with non-zero overlap with the initial state. Following [\[194\]](#), as an example, we consider the Davies map for a single qubit with $\hat{H}_s = h/2 \hat{\sigma}^z$. In [Fig. 4.3](#), we study $\Delta F(t)$ for a random

initial state (full line) and its corresponding unitarily optimised state (dashed line), which shows the occurrence of a thermal Mpemba effect $\square$. We also consider the monotones defined in Eq. (4.9) in the upper-right inset of Fig. 4.3 displaying the Mpemba crossing times as a function of α . As in the classical case discussed in Sec. 4.3.2, the variation in crossing times τ_α illustrates that the presence or absence of the Mpemba effect depends on the choice of monotone [219, 228]. Note that in Sec. 4.3.4, we demonstrate that a quantum thermal Mpemba effect can also arise in a subsystem of an initially pure state evolving under an ergodic Hamiltonian, consistent with the expectations of ETH.

4.3.4 Example 3: Thermal Mpemba Effect with ETH

In the literature, the thermal Mpemba effect has been explored mainly in Markovian open quantum systems obeying Eq. (4.21), with the recent exception of [229] where it was also found in isolated systems. Here we show that this effect can also be found in a subsystem of a quantum system evolving unitarily. A single qubit, when weakly coupled to a non-integrable environment satisfying the ETH [73, 80], relaxes to a thermal state whose temperature is determined by the total initial energy of the qubit plus environment. Following [230], we consider the total Hamiltonian

$$\hat{H}_{se} = \hat{H}_s + \hat{V} + \hat{H}_e, \quad (4.27)$$

with

$$\hat{H}_s = h_0 \hat{s}_0^z, \quad \hat{V} = \kappa \hat{s}_0^x \otimes \hat{s}_1^x, \quad (4.28)$$

where ω_0 is the system qubit's frequency and κ denotes the system-environment coupling. The environment Hamiltonian is

$$\hat{H}_e = J_z \sum_{n=1}^{N-1} \hat{s}_n^z \hat{s}_{n+1}^z + h_1 \hat{s}_1^z + h_N \hat{s}_N^z + \sum_{n=1}^N h_z \hat{s}_n^z + h_x \hat{s}_n^x. \quad (4.29)$$

We choose initial states of the form

$$|\psi_\theta\rangle_{se} = |\varphi_\theta\rangle_s \otimes |E_\theta\rangle_e \quad (4.30)$$

with $\theta = 1, 2$, $|\varphi_1\rangle = |+\rangle$ and $|\varphi_2\rangle = |0\rangle$. $|E_\theta\rangle$ is an eigenstate of $\hat{H}_e$ chosen so that the inverse temperature β of the overall state (that is fixed by the energy of $|\psi_\theta\rangle_{se}$) matches. This, together with ETH, ensures that the two different initial states on s converge to the same thermal steady state $\hat{\pi}_\beta = e^{-\beta \hat{H}_s} / \text{Tr}[-\beta \hat{H}_s]$, within the fluctuations anticipated by ETH.

The evolution of the system's state is described by a CPTP map

$$\hat{\rho}_\theta(t) = \mathcal{E}_t[\hat{\rho}_\theta] = s \text{Tr}_e[e^{-it\hat{H}_{se}}(|\varphi_\theta\rangle\langle\varphi_\theta| \otimes |E_\theta\rangle\langle E_\theta|)e^{it\hat{H}_{se}}]. \quad (4.31)$$

To quantify the divergence of $\hat{\rho}_\theta(t)$ with respect to $\hat{\pi}_\beta$, we employ the monotone of Eq. (4.10). As the map is not Markovian, $M(\hat{\rho}_\theta(t))$ is guaranteed to be monotonic only between time 0 and $t > 0$. At intermediate times, we note small fluctuations which eventually can be considered as finite size effects which would vanish completely when considering an infinitely large bath, for which the map $\mathcal{E}_t$ becomes Markovian [231]. Fig. 4.4 shows the crossing of the relative entropies indicating the occurrence of




Figure 4.4: Quantum thermal Mpemba effects in a subsystem of a system evolving unitarily. We considered the Hamiltonian Eq. (4.27) with $(J, h_z, h_x, h_1, h_N, \kappa, h_0) = (1, 0.3, 1.1, 0.25, -0.25, 0.15, 1.525)$ and $\beta = -0.46$ and $N = 15$. Inset: The crossing times τ_α obtained by studying the different α -divergence (see Eq. (4.5)).

a quantum thermal Mpemba effect $\square$. The inset shows that when using α -divergences, the crossing time τ_α depends on α .

4.4 The symmetry Mpemba effect

Besides the thermal Mpemba effect, a related phenomenon has recently drawn attention: the anomalous dynamical restoration of symmetries, first investigated in Ref. [199]. They considered unitary dynamics generated by a Hamiltonian with a continuous symmetry, where the system is initialised in a state that breaks it. For a subsystem, the symmetry is (partially) restored during the dynamics, with the remainder of the system acting as a non-Markovian environment. In this context, the Mpemba effect $\circ$ is said to occur when stronger initial symmetry breaking leads to faster symmetry restoration and the degree of asymmetry is captured by the entanglement asymmetry [199]. Instances of this effect have been observed in both integrable [232–235] and non-integrable [236–238] systems, and have been confirmed experimentally [239]. In integrable systems, a compelling physical mechanism underlying the emergence of the Mpemba effect has been identified in Ref. [232]. The key idea is that symmetry restoration is mediated by the spreading of quasiparticles that propagate with different velocities. Consequently, if the dynamics of the initially more asymmetric state involves smaller contributions from the slowest-propagating quasiparticles than those of the initially less asymmetric state, their asymmetry monotones will cross in time. This type of analysis have been applied

to the Mpemba effects for the non-Markovian evolution, with a finite-size environment. We therefore believe this applies seamlessly to our non-Markovian circuits of [Sec. 4.4.7](#).

To explain the occurrence of the Mpemba effect in this setting, we first review symmetries in (open) quantum systems in [Sec. 4.4.1](#). We then discuss the resource of asymmetry in [Sec. 4.4.2](#), and introduce modes of asymmetry, a useful tool for decomposing the dynamics in [Sec. 4.4.3](#). We show the appearance of the effect first using a classical Markov chain in [Sec. 4.4.4](#) and its quantum version in [Sec. 4.4.5](#), then a Davies map in [Sec. 4.4.6](#), and then in scenarios more closely aligned with existing literature: quantum circuits, both interacting with Markovian and with non-Markovian environments with $U(1)$ symmetry, as well as systems with non-abelian ($SU(2)$) symmetries in [Sec. 4.4.7](#).

4.4.1 Symmetries in Quantum Systems

In closed systems, Noether's theorem³ states that each symmetry implies the conservation of a corresponding physical charge [\[240, 241\]](#). When considering isolated quantum systems [\[242\]](#), the dynamics are governed by a unitary evolution operator $\hat{T} \in \mathcal{B}(\mathcal{H}_s)$. Let G be a group with unitary representation $\{\hat{U}_g\}_{g \in G} \subset \mathcal{B}(\mathcal{H}_s)$. An operator $\hat{T}$ is called G -invariant if

$$[\hat{T}, \hat{U}_g] = 0 \quad \forall g \in G, \quad (4.32)$$

or, equivalently, $\hat{T} = \hat{U}_g^\dagger \hat{T} \hat{U}_g : \forall g \in G$. In many cases of interest in physics, the group G is a continuous one-parameter family represented by unitaries of the form

$$\hat{U}_g = e^{i s_g \hat{Q}} \quad : s_g \in \mathbb{R}, \quad (4.33)$$

where s_g provides a group coordinate, while $\hat{Q}$ is a Hermitian operator, which will be conserved by all G -invariant unitary transformations, i.e., $\hat{T} \hat{Q} \hat{T}^\dagger = \hat{Q}$ for all G -invariant $\hat{T}$. This is an example of a connected Lie group. More generally, connected Lie groups of higher dimension, such as $SU(d)$, have unitary representations given by

$$\hat{U}_g = e^{i \sum_{\lambda=1}^{\Lambda} s_g^\lambda \hat{Q}_\lambda}, \quad \mathbf{s}_g^\lambda \in \mathbb{R}, \quad (4.34)$$

where $\{\hat{Q}_\lambda\}$, $\lambda = 1, \dots, \Lambda$ form a Lie algebra and are, in particular, closed under commutation

$$[\hat{Q}_a, \hat{Q}_b] = i \sum_{\lambda=1}^{\Lambda} \varepsilon_{ab}^\lambda \hat{Q}_\lambda, \quad (4.35)$$

³Noether's first theorem originally applied to continuous (Lie) groups of differentiable symmetries, but its principles have since been extended to include discrete symmetries such as parity or time reversal.

for real ε_{ab}^λ . In spin systems, four paradigmatic continuous symmetries are commonly encountered.

(i) Magnetisation conservation (phase symmetry, group $U(1)$):

$$\hat{U}_g = e^{i\gamma_g \hat{S}_z}, \quad \gamma_g \in [0, 4\pi), \quad (4.36)$$

with $\hat{S}_z = \sum_n \hat{s}_n^z$ the total spin- $\frac{1}{2}$ magnetisation about the Z -axis. In spin- $\frac{1}{2}$ chains, this $U(1)$ symmetry is formally identical to particle-number conservation in fermionic or hard-core bosonic models.

(ii) Total angular momentum conservation, associated with the rotation group $SO(3)$ or its double cover, the group $SU(2)$, is

$$\hat{U}_g = e^{i\theta_g \hat{\mathbf{n}}_g \cdot \hat{\mathbf{S}}}, \quad \theta_g \in [0, 4\pi), \quad \hat{\mathbf{n}}_g \in \mathbb{R}^3, \quad (4.37)$$

where $\hat{\mathbf{n}}_g \cdot \hat{\mathbf{n}}_g = 1$ with $\hat{\mathbf{S}} = (\hat{S}_x, \hat{S}_y, \hat{S}_z)$, $\hat{S}_\alpha = \sum_n \hat{s}_n^\alpha$, $\alpha = x, y, z$.

(iii) Energy conservation (time-translation symmetry):

$$\hat{U}_g = e^{it_g \hat{H}_s}, \quad t_g \in \mathbb{R}. \quad (4.38)$$

Here, $\hat{H}_s$ is the system Hamiltonian, and t represents time, linking this unitary operator to time-translation symmetry. In the special case where the ratios of energy levels are all rational numbers, that is, when the system's time evolution is periodic, this group of unitaries is isomorphic to $U(1)$. Otherwise, for a non-periodic system, it is isomorphic to $\mathbb{R}$ ⁴. (iv) Another important continuous symmetry, which we will not explore further, is that associated with momentum conservation (space-translation symmetry, group $\mathbb{R}^3$):

$$\hat{U}_g = e^{i\mathbf{v}_g \cdot \hat{\mathbf{P}}}, \quad \mathbf{v}_g \in \mathbb{R}^3, \quad (4.39)$$

where $\hat{\mathbf{P}} = (\hat{P}_x, \hat{P}_y, \hat{P}_z)$ is the momentum operator in the three-dimensional space, and $\hat{P}_\alpha = -i\partial_\alpha$, $\alpha = x, y, z$ in position space. Instead, we will consider its discrete version. Consider a 1D equispaced lattice with PBC on N sites, the symmetry group is $\mathbb{Z}_N$, realised by the translation operator

$$\hat{U}_g = e^{ik_g \hat{P}}, \quad k_g \in \{0, 1, \dots, N-1\}, \quad (4.40)$$

where $e^{i\hat{P}}$ is effectively the space-translation operator with unit lattice constant. Recently, the Mpemba effect has been found in this setting for qudits [243]. We note that for parity symmetry, associated with the group $\mathbb{Z}_2$, despite a thorough investigation [237, 244], the Mpemba effect has not been unambiguously observed. A

⁴Physically, the more relevant property is that the ratios of Bohr frequencies $E_i - E_j$ are rational numbers, in which case the corresponding unitary group forms a projective representation of $U(1)$.

symmetry group G acts on density matrices $\hat{\rho}$ via the unitary representation

$$\mathcal{U}_g[\hat{\rho}] = \hat{U}_g \hat{\rho} \hat{U}_g^\dagger, \quad g \in G. \quad (4.41)$$

States invariant under every $\mathcal{U}_g$ are called G -invariant. Given any state $\hat{\rho}$, its G -invariant component is extracted by the twirling map

$$\mathcal{G}[\hat{\rho}] = \int_G dg \mathcal{U}_g[\hat{\rho}] = \int_G dg \hat{U}_g \hat{\rho} \hat{U}_g^\dagger, \quad (4.42)$$

where dg denotes the normalised Haar measure on G , and we have assumed the group G is compact. This CPTP map projects any input state onto its symmetric component by averaging over all group actions. For discrete groups this becomes

$$\mathcal{G}[\hat{\rho}] = \frac{1}{|G|} \sum_{g \in G} \mathcal{U}_g[\hat{\rho}]. \quad (4.43)$$

As previously mentioned, the Mpemba effect occurs when the system interacts with an environment. If we therefore consider the G -invariant operator $\hat{T}$ of Eq. (4.32) as the evolution operator on $\mathcal{H}_s \otimes \mathcal{H}_e$, its local action on $\mathcal{H}_s$ will preserve only a weaker form of symmetry, as follows. In fact, let $\mathcal{E}$ be a map defined as

$$\mathcal{E}[\cdot] = \text{Tr}_e[\hat{T}([\cdot] \otimes \hat{\pi}_e) \hat{T}^\dagger], \quad (4.44)$$

with $\hat{\pi}_e$ a G -invariant state on $\mathcal{H}_e$. Then, it can be easily shown (see Sec. E) that, because the initial state of $\mathcal{H}_e$ and the global unitary time evolution $\hat{T}$ respect the symmetry, the time evolution of the reduced system also respects the symmetry; that is, it satisfies the G -covariance condition, also known as weak symmetry [242, 245], namely

$$\mathcal{E} \circ \mathcal{U}_g[\hat{\rho}] = \mathcal{U}_g \circ \mathcal{E}[\hat{\rho}], \quad \forall g \in G, \forall \hat{\rho}. \quad (4.45)$$

Conversely, it has been shown that any CPTP map that respects this covariance condition can be realised via Eq. (4.44), using a G -invariant unitary $\hat{T}$, and a G -invariant state $\hat{\pi}_e$ [246, 247]⁵. Therefore, while the global evolution is a G -invariant unitary transformation, the dynamics of the reduced state of the system is governed by a G -covariant CPTP map. As a result, the asymmetry (i.e., symmetry breaking) with respect to the group G can decrease. This can be quantified by means of the asymmetry monotones M , i.e. functions satisfying

$$M(\mathcal{E}[\hat{\rho}]) \leq M(\hat{\rho}) \quad \forall \hat{\rho} \in \mathcal{S}(\mathcal{H}_s) \quad (4.46)$$

for any G -covariant $\mathcal{E}$ map. Hence, while for open systems the presence of symmetry

⁵This is sometimes known as covariant Stinespring dilation theorem

does not imply conservation laws, that is, Noether's theorem is not applicable, asymmetry monotones still capture non-trivial implications of symmetries [248]. In generic open system dynamics that respect the symmetry, the asymmetry (symmetry-breaking) in the initial state decreases, although the complete restoration of the symmetry is not apriori guaranteed⁶. In the special case of closed system dynamics, the presence of symmetry implies the conservation of all asymmetry monotones. That is, for any G-invariant unitary $\hat{T}$, $M(\hat{T}\hat{\rho}\hat{T}^\dagger) = M(\hat{\rho})$: in this regime, no asymmetry can be created or destroyed. To observe Mpemba effects, thus one clearly needs to consider G-covariant maps, i.e. systems interacting with environments.

When G is the one-parameter group of time translations generated by the system Hamiltonian $\hat{H}_s$, the free operations become the set of G-covariant (also called phase-covariant) maps. An example of such maps is the TO, defined in Sec. 4.3.1. The notion of energetic coherence in quantum thermodynamics, that is, superpositions of states with different energies, can be understood as asymmetry with respect to time translations [204, 213, 251]. Note that although asymmetry under time translations is globally conserved under general energy-conserving unitaries, the reduced map can still dephase in the energy basis, effectively restoring the symmetry by erasing coherence (see Sec. 4.4.6 for further discussion). In the open systems literature, the covariance condition is often referred to as weak symmetry, to distinguish it from a stronger notion, known as strong symmetry [242, 245], in which all the Kraus operators of the channel are G-invariant (i.e., commute with the group action)⁷. Weak symmetry (G-covariance, Eq. (4.45)) admits non-increasing asymmetry monotones, whereas strong symmetry (G-invariance, Eq. (4.32)) guarantees conserved charges in the dynamics.

Under the restriction to G-covariant operations, any state exhibiting asymmetry cannot be generated for free and must instead be supplied externally as a resource. The greater the asymmetry of this resource state, the more power it confers for tasks that require breaking the symmetry. For instance, in the case of U(1) symmetry, the precision in probing the phase of a U(1) rotation increases as the state becomes more tilted with respect to the Z-direction, as quantified by Eq. (4.7). Similarly, tasks such as aligning reference frames or achieving quantum-enhanced measurements beyond the shot-noise limit benefit from using more asymmetric states [252, 253]. In this sense,

⁶A broad class of maps that are guaranteed to restore a symmetry consists of quantum Markov processes with a fixed point, as shown in Ref. [249]. However, non-Markovian examples also exist, such as spin star systems [250].

⁷For channels with multiple Kraus operators, the requirement of G-invariance (Eq. (4.32)) must hold for each Kraus element individually. Equivalently, one can demand that the Lindblad generator be unchanged under independent left- and right-unitary twists,

$$\mathcal{U}_g^{(L)} \mathcal{L} \mathcal{U}_g^{(L)\dagger} = \mathcal{U}_g^{(R)} \mathcal{L} \mathcal{U}_g^{(R)\dagger} = \mathcal{L} \quad \forall g \in G, \quad (4.47)$$

where $\mathcal{U}_g^{(L)}[\hat{\rho}] = \hat{U}_g \hat{\rho}$ (ket action) and $\mathcal{U}_g^{(R)}[\hat{\rho}] = \hat{\rho} \hat{U}_g^\dagger$ (bra action) [242, 245].

asymmetry becomes a useful resource.

4.4.2 Asymmetry as a Resource

For any symmetry group G , the resource theory of asymmetry with respect to G assumes that all systems under consideration are equipped with a representation of the group. The set of free states and free operations in this resource theory are then defined as those that respect the symmetry [210, 247, 248, 254].

In particular, for any system s with Hilbert space $\mathcal{H}_s$, the free states are the G -invariant states,

$$\mathcal{F} = \{\hat{\rho} \in \mathcal{B}(\mathcal{H}_s) \mid \mathcal{G}[\hat{\rho}] = \hat{\rho}\}, \quad (4.48)$$

where $\mathcal{G}$ denotes the corresponding twirling map.

Free operations on system s are CPTP maps $\mathcal{E}$ satisfying the covariance condition

$$\mathcal{E} \circ \mathcal{U}_g = \mathcal{U}_g \circ \mathcal{E} \quad : \forall g \in G. \quad (4.45)$$

Note that, as required in all resource theories, free operations (i.e., G -covariant maps) map free states (i.e., G -invariant states) to free states.

A useful monotone in this resource theory is the relative entropy of asymmetry [255, 256],

$$\begin{aligned} M(\hat{\rho}) &= \min_{\hat{\pi} \in \mathcal{F}} S(\hat{\rho} \parallel \hat{\pi}) \\ &= S(\hat{\rho} \parallel \mathcal{G}[\hat{\rho}]) \\ &= S(\mathcal{G}[\hat{\rho}]) - S(\hat{\rho}), \end{aligned} \quad (4.49)$$

which corresponds to Eq. (4.7). This quantifies the minimal divergence of $\hat{\rho}$ with respect to the manifold of symmetric states. Intuitively, if $\hat{\rho}$ is nearly invariant, twirling induces little entropy change; if $\hat{\rho}$ breaks the symmetry strongly, mixing over G spreads its support and increases entropy substantially [248]. The relative entropy of asymmetry can be extended to the α -divergences of Eq. (4.5) as the α -divergences of asymmetry for $\alpha \in (0, 1) \cup (1, 2]$. All these maps are guaranteed to be monotonic for G -covariant maps [248].

4.4.3 Modes of Asymmetry

Within the resource theory of asymmetry, an important tool to characterise the symmetry breaking was introduced in Ref. [204]: the modes of asymmetry. In this

section, we outline the main ingredients to this concept, a simple example, and how this idea can be used to fully understand the symmetry Mpemba effect.

Consider a finite or compact Lie group G acting unitarily on $\mathcal{H}_s$ via $\hat{U}_g : G \rightarrow \mathcal{B}(\mathcal{H}_s)$. The full operator algebra $\mathcal{B}(\mathcal{H}_s)$ then carries an induced adjoint action of G , and the irreducible representation of G labelled by μ , which appears with multiplicity m_μ in $\mathcal{H}_s$. One can use the irreducible tensor operator basis $\{\hat{T}_\alpha^{(\mu,q)}\}$ of $\mathcal{B}(\mathcal{H}_s)$ $q = 1, \dots, d_\mu$ (d_μ dimension of the irrep, namely an irreducible representation) and $\alpha = 1, \dots, m_\mu$ satisfying $\text{Tr}[\hat{T}_\alpha^{(\mu,q)\dagger} \hat{T}_\beta^{(\nu,\ell)}] = \delta_{\mu\nu} \delta_{q\ell} \delta_{\alpha\beta}$, to decompose any operator $\hat{X} \in \mathcal{B}(\mathcal{H}_s)$ as

$$\hat{X} = \sum_{\mu,q} \hat{X}^{(\mu,q)}. \quad (4.50)$$

Here, $\hat{X}^{(\mu,q)}$ is the (μ, q) -mode of $\hat{X}$, defined as

$$\hat{X}^{(\mu,q)} = \sum_{\alpha} \text{Tr}[\hat{T}_\alpha^{(\mu,q)\dagger} \hat{X}] \hat{T}_\alpha^{(\mu,q)}. \quad (4.51)$$

A basis independent way to define these modes is by the group's action

$$\hat{X}^{(\mu,q)} = d_\mu \int_G dg D_{qq}^{(\mu)*}(g) \mathcal{U}_g[\hat{X}], \quad (4.52)$$

where $D_{\ell q}^{(\mu)}(g) = \langle \mu, \ell | \hat{U}_g | \mu, q \rangle$. We write $\mathcal{H}^{(\mu,q)}$ for the subspace where the (μ, q) -mode lives. Assuming $\mu = 0$ corresponds to the trivial irrep, with $d_\mu = 1$, one has $D_{00}^{(0)}(g) = 1$ from which it becomes evident how the mode associated with the trivial irrep corresponds to the twirled part of the operator, as in Eq. (4.42). By orthonormality of the $\{\hat{T}_\alpha^{(\mu,q)}\}$ basis it also follows that all the other modes will characterise a symmetry breaking. For discrete groups Eq. (4.52) becomes

$$\hat{X}^{(\mu,q)} = \frac{d_\mu}{|G|} \sum_{g \in G} D_{qq}^{(\mu)*}(g) \mathcal{U}_g[\hat{X}]. \quad (4.53)$$


This formalism becomes particularly useful when the system evolves by the free operations of the resource theory of asymmetry. Specifically, following Ref. [204], a map $\mathcal{E}$ is G -covariant iff it acts within each μ -sector as

$$\mathcal{E}[\hat{X}^{(\mu,q)}] = (\mathcal{E}[\hat{X}])^{(\mu,q)}. \quad (4.54)$$

Namely, each mode evolves independently and so, since the asymmetry can only decrease by G -covariant maps and the trace norm respects the data processing inequality, the trace norm of each asymmetry mode is individually a monotone, that is

$$\|\mathcal{E}[\hat{X}^{(\mu,q)}]\|_1 \leq \|\hat{X}^{(\mu,q)}\|_1. \quad (4.55)$$

Figure 4.5: U(1) symmetry sectors and mode occupancy for three-qubit tilted states. Basis operators $|m, \alpha_1\rangle\langle m', \alpha_2|$ carry magnetisation difference $\mu = m - m'$. Ordering rows/columns by increasing m groups the 8×8 operator basis into fixed- $|\mu|$ blocks (a). (b–d): mode-occupancy histograms Eq. (4.61) for the tilted-ferromagnet family Eq. (4.60) with $N_s = 3$, from symmetric ($\mu = 0$ only) to maximally asymmetric (uniform across sectors).



As a concrete example, consider the symmetry group $G = \text{U}(1)$, generated by the total- z magnetisation operator $\hat{S}_z$ on N_s spin-half particles. Its unitary representation is as in Eq. (4.36) given by $\hat{U}_\gamma$ with $\gamma \in [0, 4\pi)$.

Because of the group's simple structure, its asymmetry modes are labelled by a single integer μ , which determines the U(1) irrep as $D^{(\mu)} : e^{i\theta} \mapsto e^{i\mu\theta}$. The operator $\hat{S}_z$ is diagonal in the computational basis $\{|0\rangle, |1\rangle\}^{\otimes N_s}$, with eigenvalues $\pm\frac{1}{2}, \dots, \pm N_s/2$ when N_s , the number of qubits, is odd, and $0, \pm 1, \dots, \pm N_s/2$ when N_s is even. We can equivalently label the elements of the computational basis as $\{|m, \alpha\rangle\}$, where m is the eigenvalue of $\hat{S}_z$, and α is a multiplicity index. In particular, m is simply half the difference between the number of 0's and 1's in a computational basis element.

Then, the μ -component of any $\hat{\rho}$ is given by projection

$$\begin{aligned} \hat{\rho}^{(\mu)} &= \int_0^{4\pi} \frac{d\gamma}{4\pi} e^{-i\gamma\mu/2} (e^{i\gamma\hat{S}_z} \hat{\rho} e^{-i\gamma\hat{S}_z}) \\ &= \sum_m \hat{\Pi}_m \hat{\rho} \hat{\Pi}_{m+\mu}, \end{aligned} \quad (4.56)$$

where $\hat{\Pi}_m = \sum_\alpha |m, \alpha\rangle\langle m, \alpha|$ is the projection to the subspace of states with m -magnetisation. This can also be written as Eq. (4.51) in terms of the irreducible tensor basis of U(1) as

$$\hat{\rho}^{(\mu)} = \sum_\alpha \text{Tr}[\hat{T}_\alpha^{(\mu)\dagger} \hat{\rho}] \hat{T}_\alpha^{(\mu)}, \quad (4.57)$$

where one can choose the tensor operators to be simply the outer product of two computational basis elements, or, equivalently, two eigenvectors of $\hat{S}_z$.

$$\hat{T}_\alpha^{(\mu)} = |m, \alpha_1\rangle\langle m + \mu, \alpha_2| \quad (4.58)$$

where $\alpha = (m, \alpha_1, \alpha_2)$. It follows that

$$e^{i\gamma\hat{S}_z} \hat{\rho}^{(\mu)} e^{-i\gamma\hat{S}_z} = e^{i\gamma\mu} \hat{\rho}^{(\mu)}, \quad (4.59)$$

where at $\mu = 0$ we find the symmetrised (twirled) state.

Fixing the number of qubits to $N_s = 3$, we order the basis states by increasing m_k (see panel (a) of Fig. 4.5) to find the block structure with $\mu = 0, \pm 1, \pm 2, \pm 3$. To illustrate the mode decomposition, we employ the family of tilted pure states

$$|\varphi(\theta)\rangle = \bigotimes_{n=1}^3 e^{-i\theta s_n^y} |0\rangle_n, \quad (4.60)$$

so that $\theta = 0$ gives the fully Z -aligned (symmetric) state, $\theta = \pi/4$ an intermediate tilt, and $\theta = \pi/2$ the fully x -aligned (maximally asymmetric) state (panels (b-d) of Fig. 4.5). Writing $\hat{\rho}_\theta = |\varphi(\theta)\rangle\langle\varphi(\theta)|$, we measure the weight of mode μ via the rescaled trace norm,

$$p_\theta^{(\mu)} = \|\hat{\rho}_\theta^{(\mu)}\|_1 / \|\hat{\rho}_{\pi/2}^{(\mu)}\|_1. \quad (4.61)$$

As θ increases from 0 to $\pi/2$, the weights $p_\theta^{(\mu \neq 0)}$ grow, showing that the tilted state acquires higher- μ asymmetry components. As we discuss in the next section, this formalism helps unify the two phenomena: the thermal Mpemba effect hinges on the small overlap with the overall slowest-decaying eigenmode, whereas the symmetry Mpemba effect stems exclusively from the decay rate of the slowest eigenmode living in a ($\mu \neq 0$)-sector.

4.4.4 Example 1: A Classical Symmetry Mpemba Effect with $\mathbb{Z}_4$

To illustrate how an initially more asymmetric distribution can restore symmetry faster than a more symmetric one , we consider the minimal classical setup in which distinct symmetry-breaking eigenmodes (namely eigenmodes living in $\mu \neq 0$ -sectors) decay at different rates: a four-site Markov chain on a ring (represented in Fig. 4.6), invariant under the cyclic rotation $n \mapsto n + 1 \pmod{4}$, namely $\mathbb{Z}_4$ -covariant. The irreps are $D^{(\mu)}(n) = e^{2\pi i n \mu / 4} = i^{n\mu}$, where $\mu = 0, 1, 2, 3$. Jumps occur at a unit rate between each pair of neighbours, as described by the Liouvillian

$$\hat{\mathcal{L}}_{nm}^{\text{cl}} = \begin{cases} 1 + \varepsilon, & m = n + 1 \pmod{4}, \\ 1 - \varepsilon, & m = n - 1 \pmod{4}, \\ -2, & n = m, \\ 0, & \text{otherwise,} \end{cases} \quad (4.62)$$

where $\varepsilon \in (-1, 1)$ is the chirality parameter. In this case, there are four modes of asymmetry, corresponding to $\mu = 0, 1, 2, 3$. Then, the analogue of tensor operators are given by

$$\mathbf{T}^{(\mu)} = \frac{1}{2}(1, i^\mu, i^{2\mu}, i^{3\mu})^\top. \quad (4.63)$$


In particular, under the action of cyclic rotation $n \mapsto n + 1 \pmod{4}$, this vector is mapped to itself, up to a phase $-i^\mu$. Writing the eigenmodes of $\hat{\mathcal{L}}^{\text{cl}}$ in terms of the $\{\mathbf{T}^{(\mu)}\}$ basis we get $\boldsymbol{\pi} = \mathbf{T}^{(0)}$ (eigenvalue $\lambda_1 = 0$), $\mathbf{r}_2 = \mathbf{T}^{(1)}$ ($\lambda_2 = -2 + 2i\varepsilon$), $\mathbf{r}_3 = \mathbf{T}^{(3)}$ ($\lambda_3 = -2 - 2i\varepsilon$), $\mathbf{r}_4 = \mathbf{T}^{(2)}$ ($\lambda_4 = -4$).

Hence, for $\varepsilon \in (-1, 1)$, $\mathbf{r}_4$ is the fastest decaying mode. (Note that the detailed balance condition is violated for $\varepsilon \neq 0$.) Each eigenmode lives in a single asymmetry sector, which guarantees the preservation of the mode structure, as in Eq. (4.54). Only $\boldsymbol{\pi}$ (which lives in $\mu = 0$) is fully symmetric; the three asymmetric eigenmodes -two slow ones ($\mathbf{r}_2, \mathbf{r}_3$) and one fast one ($\mathbf{r}_4$)- govern how different distortions decay back to equilibrium. In this case, the relative entropy of asymmetry in Eq. (4.49) is equal to the KL-divergence of the current distribution $\mathbf{p}(t)$ from its symmetrised version, $\hat{\mathcal{G}}\mathbf{p}(t)$, adapted to this case using the discrete version of the twirling operation in Eq. (4.43). Since averaging over all four rotations yields the uniform state $\hat{\mathcal{G}}\mathbf{p}(t) = \boldsymbol{\pi}/2$, the relative entropy of asymmetry simplifies to

$$M(\mathbf{p}(t)) = D(\mathbf{p}(t) \parallel \hat{\mathcal{G}}\mathbf{p}(t)) = \sum_{k=1}^4 p_k(t) \ln(4 p_k(t)), \quad (4.64)$$

which vanishes precisely when $\mathbf{p}$ is uniform and decreases under the dynamics $e^{t\hat{\mathcal{L}}^{\text{cl}}}$. To find a strong Mpemba effect, we compare two probability vectors chosen along

Figure 4.6: (a) the generator of the classical Markov chain with the four sites arranged on a ring and three type of jumps making them interacts. (b) and (c) the distributions $\mathbf{p}_1$ and $\mathbf{p}_2$ which are $(0.5, 0, 0.5, 0)^\top$ and $(0.5, 0.25, 0, 0.25)^\top$ respectively.



orthogonal spectral directions as in Eq. (4.18):

$$\mathbf{p}_1 = \underbrace{\frac{1}{2}\boldsymbol{\pi}}_{\mathbf{p}_1^{(0)}} + \underbrace{\frac{1}{4}\mathbf{r}_2}_{\mathbf{p}_1^{(1)}} + \underbrace{\frac{1}{4}\mathbf{r}_3}_{\mathbf{p}_1^{(3)}}, \quad \mathbf{p}_2 = \underbrace{\frac{1}{2}\boldsymbol{\pi}}_{\mathbf{p}_2^{(0)}} + \underbrace{\frac{1}{2}\mathbf{r}_4}_{\mathbf{p}_2^{(2)}}, \quad (4.65)$$

where the first lies entirely along the slow symmetry-breaking eigenmodes $\mathbf{r}_2$ and $\mathbf{r}_3$, while the second has only a component in the fast eigenmode $\mathbf{r}_4$. The underbraces highlight the mode decomposition of the states where $\mathbf{p}_\theta^{(\mu)} = \langle \mathbf{T}^{(\mu)}, \mathbf{p}_\theta \rangle_\pi \mathbf{T}^{(\mu)}$ for $\theta = 0, 1$ which corresponds to the adaptation of Eq. (4.51) to this case. While at $t = 0$, $0.35 \approx M(\mathbf{p}_1) < M(\mathbf{p}_2) \approx 0.69$, if we pick $\varepsilon = 1/4$, according to Eq. (4.19) their asymmetry behaviour asymptotically decreases as

$$M(\mathbf{p}_1(t)) \sim \frac{1}{4}e^{-4t}, \quad M(\mathbf{p}_2(t)) \sim \frac{1}{2}e^{-8t}. \quad (4.66)$$

By numerically propagating the two initial configurations Eq. (4.65) with the Liouvillian

Eq. (4.62), we find that the two monotones M cross at time $\tau = 0.16$. In other words, the distribution $\mathbf{p}_2$ outpaces $\mathbf{p}_1$ in restoring the $\mathbb{Z}_4$ symmetry, manifesting the symmetry Mpemba effect in a classical system. In this minimal example, the eigenmodes of the Liouvillian coincide with the irreducible basis, and the uniqueness of the symmetric configuration reduces symmetry restoration to ordinary equilibration on this classical Markov chain. This simplification, however, is an artefact of our minimal model. In the next section, we turn to a richer scenario in which each eigenvalue sector μ appears with multiplicity four. In that case, one can observe genuinely new behaviour, such as nontrivial evolution within the $\mu = 0$ symmetric subspace, that cleanly separates the process of symmetry restoration from the usual equilibration dynamics.

4.4.5 Example 2: A Quantum Symmetry Mpemba Effect with $\mathbb{Z}_4$

In the classical Markov chain of Sec. 4.4.4, the generator splits into four one-dimensional symmetry sectors under the action of $\mathbb{Z}_4$. We now construct a quantum analogue: a single spinless particle hopping on a four-site ring, governed by a $\mathbb{Z}_4$ -covariant Lindblad dynamics. The irreducible tensor operators are labelled by $\mu \in \{0, 1, 2, 3\}$ with multiplicity index $\alpha \in \{0, 1, 2, 3\}$:

$$\hat{T}_\alpha^{(\mu)} = \frac{1}{2} \sum_{n=1}^4 e^{2\pi i n \mu / 4} |n + \alpha \pmod{4}\rangle \langle n|. \quad (4.67)$$

Here n labels the site on the ring, understood modulo 4 so that site 5 is identified with site 1 and site 0 with site 4, and the phase factor ensures each $\hat{T}_\alpha^{(\mu)}$ transforms in the μ -th irreducible representation of $\mathbb{Z}_4$. To define the dynamics, we embed the classical jump-rates into Lindblad hopping operators and add a coherent Hamiltonian term. The master equation reads

$$\frac{d\hat{\rho}}{dt} = \mathcal{L}[\hat{\rho}] = -i[\hat{H}_s, \hat{\rho}] + \mathcal{D}_-[\hat{\rho}] + \mathcal{D}_+[\hat{\rho}], \quad (4.68)$$

where the coherent part of the dynamics is generated by the Hamiltonian

$$\hat{H}_s = J \sum_{n=1}^4 |n\rangle \langle n+1 \pmod{4}| + |n+1 \pmod{4}\rangle \langle n| \quad (4.69)$$

which describes hopping between neighbouring sites on a four-site ring with amplitude J . The incoherent, biased hopping around the ring is implemented by the two dissipators


(defined as Eq. (4.22)) respectively for the jump operators

$$\begin{aligned}\hat{L}_{n,+} &= \sqrt{1+\varepsilon} \left| n+1 \pmod{4} \right\rangle\langle n|, \\ \hat{L}_{n,-} &= \sqrt{1-\varepsilon} \left| n-1 \pmod{4} \right\rangle\langle n|,\end{aligned}\tag{4.70}$$

where the bias parameter $\varepsilon \in [-1, 1]$, as before, tilts the rates so that $\varepsilon > 0$ enhances clockwise jumps and $\varepsilon < 0$ enhances counterclockwise jumps. This construction yields a Lindbladian $\mathcal{L}$ that is manifestly $\mathbb{Z}_4$ -covariant and thus block-diagonal in the μ -sectors.

We set $J = 1$ and $\varepsilon = 1/4$. The Lindblad superoperator $\mathcal{L}$ admits sixteen eigenmodes, four for each μ -sector, which can be expressed in the tensor basis $\hat{T}_\alpha^{(\mu)}$. Among these, the unique eigenmode belonging to the $\mu=0$ -sector and zero eigenvalue is the steady state $\hat{\pi}$. The remaining fifteen eigenmodes $\hat{r}_k$ decay with eigenvalues λ_k , ordered by decreasing real part $\text{Re } \lambda_k$ ⁸.

Figure 4.7: Symmetry-breaking quantum Mpemba effect for a $\mathbb{Z}_4$ -covariant dynamics. Main panel: relative entropy of asymmetry $M(\hat{\rho}_{1,2}(t))$, decaying exponentially with rates set by λ_{10} and λ_{16} . Inset: sector trace norms $\|\hat{\rho}_{1,2}^{(\mu)}(t)\|_1$, illustrating that eigenmodes $\hat{r}_{10}$ and $\hat{r}_{14}$ are related by Hermitian conjugation and overlap in the $\mu = 1$ and $\mu = 3$ sectors respectively; the structure of $\hat{\rho}_1(t)$ induces an oscillatory decay of $\|\hat{\rho}_1^{(2)}(t)\|_1$.



We prepare two initial states by mixing the steady state with selected decaying eigenmodes:

$$\begin{aligned}\hat{\rho}_1 &= \underbrace{\hat{\pi}/2 + (\hat{r}_2 - \hat{r}_3 + \hat{r}_4)/9}_{\hat{\rho}_1^{(0)}} + \underbrace{(-\hat{r}_5/5)}_{\hat{\rho}_1^{(1)}} + \underbrace{(\hat{r}_{14} + \hat{r}_{15})/8}_{\hat{\rho}_1^{(2)}} + \underbrace{(-\hat{r}_{13}/5)}_{\hat{\rho}_1^{(3)}}, \\ \hat{\rho}_2 &= \underbrace{\hat{\pi}/2 + (\hat{r}_2 - 2\hat{r}_3 + \hat{r}_4)/10}_{\hat{\rho}_2^{(0)}} + \underbrace{-\hat{r}_{16}/2}_{\hat{\rho}_2^{(2)}}.\end{aligned}\tag{4.71}$$

These choices introduce different amounts of initial asymmetry in sectors $\mu = 1, 2, 3$. We then evolve each state under the dynamics

$$\hat{\rho}_\theta(t) = e^{t\mathcal{L}}[\hat{\rho}_\theta],\tag{4.72}$$

and quantify the restoration of $\mathbb{Z}_4$ symmetry via the relative entropy of asymmetry $M(\hat{\rho}_\theta(t))$ as defined in Eq. (4.49). Fig. 4.7 shows that the state $\hat{\rho}_2$, which starts with

⁸In the case of equal real parts, modes are further ordered by increasing μ , and then by decreasing imaginary part.

a larger asymmetry in sector $\mu = 2$, relaxes faster than $\hat{\rho}_1$, which has its dominant asymmetry in sector $\mu = 1$ (and equivalently $\mu = 3$). Despite having greater initial asymmetry, $\hat{\rho}_2$ crosses below $\hat{\rho}_1$ in $M(t)$, demonstrating the symmetry Mpemba effect. The inset highlights the decay of each charge-sector norm, confirming the underlying exponential behaviour of the asymmetry modes. In this example we see a first advantage brought by the modes of asymmetry: they help identify the eigenmodes responsible for the symmetry breaking.

4.4.6 Example 3: Time-Translation-Symmetry Mpemba Effect in Davies Maps

We now turn to a continuous symmetry, namely the time-translation symmetry. Let us consider once again a Davies map described in Sec. 4.3.3 but now focusing on the symmetry Mpemba effect . When the symmetry group action $e^{i\hat{H}_s t} : t \in \mathbb{R}$ is generated by the system Hamiltonian $\hat{H}_s$, the modes of asymmetry $\hat{\rho}^{(\mu)}$ of a state $\hat{\rho}$ coincide with its energy coherence blocks at Bohr frequency μ . Concretely, $\hat{\rho} = \sum_{\mu} \hat{\rho}^{(\mu)}$ and $\mathcal{U}_g[\hat{\rho}^{(\mu)}] = e^{i\mu t_g} \hat{\rho}^{(\mu)}$, where $\mathcal{U}_g[\cdot] = e^{i\hat{H}_s t_g}[\cdot]e^{-i\hat{H}_s t_g}$, so that $\mu = 0$ picks out all diagonal (incoherent) terms and each $\mu \neq 0$ selects the off-diagonal coherences oscillating at frequency μ under unitary time evolution. The Lindbladian eigenmodes corresponding to symmetry-preserving (symmetry-breaking) modes are depicted in green (orange) in the lower-left inset of Fig. 4.8. In Davies maps, the Lindbladian $\mathcal{L}$ commutes by construction with the Hamiltonian evolution and thus respects the time-translation symmetry [216]. (Recall that when, up to a shift, the energy levels are all integer multiples of a fixed energy, i.e., the system dynamics is periodic, time translation becomes isomorphic to $U(1)$, and such maps are also referred to as phase-covariant.) As discussed in Sec. 4.4.1, this guarantees the monotonicity of the functions $M_{\alpha}(\hat{\rho}(t))$. A symmetry Mpemba effect  occurs when the state with higher initial coherence decoheres faster than one with lower coherence, and their monotone curves cross in time.

As a concrete illustration, consider the transverse-field Ising chain

$$\hat{H}_s = -J \sum_{n=1}^{N_s-1} \hat{\sigma}_n^z \hat{\sigma}_{n+1}^z + h \sum_{n=1}^{N_s} \hat{\sigma}_n^x, \quad (4.73)$$


with $N_s = 4$ and OBCs. We compare two unitarily connected initial states

$$\hat{\rho}_1 = (\hat{\rho}_{\text{Gibbs}} + \gamma \hat{\rho}_{\text{rand}})/N, \quad (4.74)$$

$$\hat{\rho}_2 = \hat{R}_{\text{metr}} \hat{\rho}_1 \hat{R}_{\text{metr}}^{\dagger}, \quad (4.75)$$

where N is a normalisation constant, $\hat{\rho}_{\text{rand}} = \hat{X}^{\dagger} \hat{X} / \text{Tr}[\hat{X}]$, with $\hat{X}$ being a random

Figure 4.8: Symmetry-breaking Mpemba effect for a transverse-field Ising model (TFIM) Eq. (4.73) thermalising under a Davies map [205]. Parameters: $J = h = 1$, $N_s = 4$, $\beta = 0.1$, $\beta_i = 1$, $\gamma = 0.05$ (no left/right hopping asymmetry, $\varepsilon = 0$). Solid: $M(\hat{\rho}_1(t))$ with $\hat{\rho}_1 = (\hat{\rho}_{\text{Gibbs}} + \gamma\hat{\rho}_{\text{rand}})/N$; dashed: $M(\hat{\rho}_2(t))$ with $\hat{\rho}_2 = \hat{R}_{\text{metr}}\hat{\rho}_1\hat{R}_{\text{metr}}^\dagger$, where $\hat{R}_{\text{metr}}$ (Metropolis [194]) minimises overlap with the slow mode $\hat{\ell}_2$. Insets: (a) crossing time vs α for Petz-Rényi divergences Eq. (4.5); (b) ten slowest eigenvalues of the Davies generator, with $\mu = 0$ (green) vs $\mu \neq 0$ (orange).



Hermitian matrix with real and imaginary components uniformly distributed in $[0, 1]$ and $\hat{R}_{\text{metr}}$ is the unitary transformation minimising the overlap of $\hat{\rho}_1$ with $\hat{\ell}_2$ via the Metropolis algorithm [194]. Fig. 4.8 shows the crossing of $M(\hat{\rho}_1(t))$ (solid) and $M(\hat{\rho}_2(t))$ (dashed). Moreover, the upper-right inset indicates that when considering the Petz-Rényi α -divergences, the crossing time decreases upon increasing α . Importantly, in contrast to the thermal Mpemba effect (Sec. 4.3.3), that involves both $\mu = 0$ (populations) and $|\mu| > 0$ (coherences) blocks, here only the $|\mu| > 0$ modes contribute to the decay of $M(\hat{\rho}_1(t))$: the slowest-decaying coherent mode corresponds to the complex eigenvalue with real part closest to zero.

4.4.7 Example 4: Symmetry Mpemba Effects in Quantum circuits

Originally, the symmetry Mpemba effect was investigated in continuous-time models [199]. Here we consider circuit models [236, 237], in which time is treated as a discrete variable, and focus on two symmetries, $U(1)$ and $SU(2)$. We consider N qubits, with N even, where the first N_s qubits define the system and the remaining N_e the environment. The initial state spontaneously breaks the symmetry on the system qubits only, with the extent of the breaking controlled by a parameter θ . The state is then evolved by a circuit composed of symmetry-preserving gates arranged in a brickwork geometry. We consider two-qubit gates and the Floquet unitary $\hat{U}$ is

$$\hat{U} = \prod_{n=1}^{N/2} \hat{u}_{2n-1} \prod_{n=1}^{N/2} \hat{u}_{2n} , \quad (4.76)$$

where $\hat{u}_n$ are gates on the $n, n+1$ qubits and $\hat{u}_N$ connects the last and the first qubits to enforce periodic boundary conditions, as represented in the two-qubit brickwork circuits of Fig. 4.9(b). While in the absence of symmetries any unitary can be decomposed into two local gates, when enforcing a symmetry this is not guaranteed anymore: a

generic G-invariant unitary might not be composed of two-local G-invariant gates [257, 258]. (For U(1) and SU(2), two-local gates suffice for semi-universality, meaning that for states restricted to a single irrep of the symmetry, the locality of gates imposes no additional constraints.) We initialise the system and environment as

$$(\hat{\sigma}_\theta)_{se} = (\hat{\rho}_\theta)_s \otimes \hat{\pi}_e, \quad (4.77)$$

picking states that break the symmetry only on the system qubits, and controlling the degree of symmetry breaking through a tilt parameter θ , while $\hat{\pi}_e$ is a G-invariant state. This guarantees that the system evolution map

$$\mathcal{E}[\hat{\rho}_\theta] = \text{Tr}_e[\hat{U} \hat{\sigma}_\theta \hat{U}^\dagger] \quad (4.78)$$

is a non-unitary CPTP map that is G-covariant; namely, $\mathcal{E}$ is a free operation for the asymmetry resource (see Sec. E). At each time step, we compute the residual




Figure 4.9: Circuit schematics for Sec. 4.4.7. (a): initialisation, with system qubits (pink) in an asymmetric state set by θ and environment qubits (blue) in a symmetric state. (b): brickwork evolution $\hat{U}$ built from random two-qubit gates constrained to U(1) or SU(2) symmetry. (b1): Markovian channel via environment reset to the maximally mixed (identity) state after each $\hat{U}$. (b2): non-Markovian channel with no mid-circuit reset/measurement, allowing information backflow from the environment.

asymmetry with the relative entropy of asymmetry $M(\hat{\rho}_\theta(t))$, as in Eq. (4.49).




Figure 4.10: Symmetry-breaking Mpemba effect in a Markovian U(1) model averaged over 100 circuit realisations with $N_s = 4$, $N_e = 2$. Panel (a) shows the time evolution of the monotones. The inset, panel (b) instead displays the averaged spectrum of the Markovian map, showing how the slow eigenmodes tend to come from low $|\mu|$ -sectors, explaining the strong Mpemba effect $\circ$ obtained by choosing these three states. Mean (dots) and spread (shaded) of the real part of the slowest decay rate in each $|\mu|$ -sector, showing a clear, monotonic decrease with $|\mu|$. It is evident that as $|\mu|$ increases, $\text{Re} \lambda_\mu$ decreases monotonically. (d) Overlap $\text{Tr}[\hat{\ell}_\mu^\dagger \hat{\rho}_{\theta,b}]$ of each initial state with its sector's slowest eigenmode $\hat{\ell}_\mu$, confirming that each state couples only to sectors that are multiples of b .

Symmetry: U(1)

We now explore how U(1) symmetry (i.e. conservation of total $\hat{S}^z$) can give rise to a Mpemba effect in random circuits. In particular, we show how it is possible to achieve a strong Mpemba effect. Considering the exact model discussed in Ref. [2] and Sec. 2.6, the two-qubit gates are

$$\hat{u}_{n,n+1} = e^{-i\hat{h}_{n,n+1}} e^{-ih_n\hat{s}_n^z} e^{-ih'_n\hat{s}_{n+1}^z}, \quad (4.79)$$

with an XXZ-type interaction

$$\hat{h}_{n,n+1} = \frac{J}{2} \left(\hat{s}_n^+ \hat{s}_{n+1}^- e^{i\phi_n} + \hat{s}_n^- \hat{s}_{n+1}^+ e^{-i\phi_n} \right) + J_z \hat{s}_n^z \hat{s}_{n+1}^z, \quad (4.80)$$

where $e^{-ih_n\hat{s}_n^z}$ and $e^{-ih'_n\hat{s}_{n+1}^z}$ are local Z-rotations. Here $\hat{s}_n^\pm = \hat{s}_n^x \pm i\hat{s}_n^y$, and ϕ_n is the Peierls phase. By sampling the five parameters $h_n, h'_n, \phi_n, J, J_z$ independently for each gate, the full circuit $\hat{U}$ remains U(1)-invariant, conserving total magnetisation $\hat{S}^z$.

To study the Mpemba effect in circuits we first discuss the clean Markovian limit with a memoryless environment, and then generalise to non-Markovian circuits with a finite number of environment qubits. In the Markovian limit, the environment is reset after each layer, so that

$$\mathcal{E}_t = (\mathcal{E})^t, \quad \mathcal{E} = \mathcal{E}_1. \quad (4.81)$$

The environment state $\hat{\pi}_e$ is the maximally mixed state (see Fig. 4.9(b)), each map $\mathcal{E}$ is U(1)-covariant (which guarantees monotonicity of $M(\hat{\rho}(t))$) and thus decomposes into asymmetry sectors

$$\mathcal{E} = \bigoplus_{\mu} \mathcal{E}^{(\mu)} \quad (4.82)$$

where $\mathcal{E}^{(\mu)}$ acts on the subspace of the operator space associated with mode μ . We can diagonalise each $\mathcal{E}^{(\mu)}$ individually using the basis $\{\hat{T}_\alpha^{(\mu)}\}$ as defined in Eq. (4.57) and obtain the decay eigenmodes in each μ -sector.

To produce a strong Mpemba effect, we prepare product states of N_s qubits in blocks of size b with a uniform tilt θ :

$$|\varphi(\theta, b)\rangle = \bigotimes_{n=1}^{\lceil N_s/b \rceil} \exp\left[-i\theta \prod_{j=1}^b \hat{s}_{(n-1)b+j}^y\right] |0\rangle^{\otimes N_s}. \quad (4.83)$$

By varying θ and b , we select system states $\hat{\rho}_{\theta,b} = |\varphi(\theta, b)\rangle\langle\varphi(\theta, b)|$ whose support lies in different μ -sectors (e.g. $b = 1$ populates all μ -sectors, $b = 2$ populates only the even μ -sectors, $b = 3$ only sectors multiple of 3, namely $\hat{\rho}_{\theta,b} = \sum_{k=0}^{\lceil N_s/b \rceil - 1} \hat{\rho}_{\theta,b}^{(kb)}$), so each state

$\hat{\rho}_{\theta,b}$ overlaps differently with the slowest eigenmodes $\{\hat{\ell}_\mu\}_\mu$ (as in Fig. 4.10(d)) where

$$\hat{\ell}_\mu = \hat{\ell}_{k(\mu)}, \quad \lambda_\mu = \lambda_{k(\mu)}, \quad (4.84)$$

where $k(\mu)$ picks out the slowest eigenmode in the μ -sector

$$k(\mu) = \arg \max_{k: \hat{\ell}_k \in \mathcal{H}_s^{(\mu)}} \text{Re } \lambda_k, \quad (4.85)$$

where $\mathcal{H}_s^{(\mu)}$ is the μ -sector of $\mathcal{H}_s$. The main panel of Fig. 4.10(a) then reveals distinct asymptotic decay exponents, i.e. a strong Mpemba effect. Avoiding overlap with specific μ -sectors induces a strong Mpemba effect because, in almost every realisation of $\hat{U}$, sectors with larger $|\mu|$ decay more quickly: for $|\mu| > |\mu'|$, one finds $\text{Re } \lambda_\mu < \text{Re } \lambda_{\mu'}$. Indeed, as illustrated in Fig. 4.10, the slowest eigenmodes reside in the blocks with the smallest $|\mu|$, that is, the lowest absolute U(1) charge. While a more thorough explanation of this phenomenon requires further investigation, one can consider the following possible mechanisms: (i) Modes with higher U(1) charge tend to evolve more rapidly, leading to faster build-up of correlations with the environment and, consequently, quicker information leakage and decay. (ii) Since $\mathcal{E}$ is constructed from random U(1)-conserving gates, the blocks corresponding to smaller $|\mu|$, being larger in dimension, are statistically more likely to contain the slowest-decaying eigenmodes. Moreover, just as in the example of Sec. 4.4.6, the initial state's overlap with the slow eigenmodes in sectors with $\mu \neq 0$ governs its relaxation rate. Note that the relative entropy of asymmetry $M(\hat{\rho}(t))$ cannot be interpreted as entanglement asymmetry in this scenario, as the global state is not pure.




Figure 4.11: Symmetry-breaking Mpemba effect in a non-Markovian U(1) model averaged over 100 circuit realisations with $N_s = 8$, $N_e = 12$. The time evolution of the monotones reveals how the more tilted state, although it starts at a higher amount of asymmetry, restores the symmetry faster. Different from the Markovian case, the monotones do not decay either monotonically or exponentially, as the asymmetry is just distributed in the whole system + environment and therefore finite size effects become more evident.

While the Markovian case allows for a clean explanation of the effect, most studies of the symmetry Mpemba effect use non-Markovian circuits. To connect to these, we again prepare a tilted product state on the system (with block size $b = 1$) with

$\hat{\rho}_\theta = |\varphi(\theta)\rangle\langle\varphi(\theta)|$ and

$$|\varphi(\theta)\rangle = \bigotimes_{n=1}^{N_s} e^{-i\theta \hat{s}_n^y} |0\rangle_n \quad (4.86)$$

matching Refs. [199, 237] except that the tilt now spans only the N_s system qubits, while the environment remains symmetric. Indeed, we initialise the environment in the computational zero state $\hat{\pi}_e = |0\rangle\langle 0|_e = |m_{\min}\rangle\langle m_{\min}|_e$, where $|m\rangle$ denotes the equal superposition of all basis states with magnetisation m and $\hat{s}_n^\alpha$ are the spin- $\frac{1}{2}$ operators (with $\alpha \in \{x, y, z\}$). Eq. (4.82), which in the Markovian case describes how each symmetry-sector mode decays at a fixed rate, still applies when the map is non-Markovian. However, now it is time-dependent, so both the eigenmodes and the stationary state can shift at each step, while the eigenvalues usually shrink exponentially with time when $\hat{U}$ is ergodic [259]. As already discussed, blocks with larger $|\mu|$ admit fewer eigenvalues near unity, so their modes decay more rapidly. Consequently, an initially highly asymmetric state, which has a larger weight in those fast-decaying, high-frequency sectors, will shed its asymmetry more quickly than a less asymmetric one, yielding the Mpemba effect.

Fig. 4.11 tracks five such states (different θ) on a system of $N_s = 8$ qubits coupled to $N_e = 12$ environment qubits.

Under general covariant time evolutions, asymmetry measures at time $t > 0$ are bounded by their value at $t = 0$, although in generic non-Markovian processes they can still exhibit transient increases at intermediate times. The deviations observed in Fig. 4.11 are therefore a consequence of information backflow [260] between the system and environment, amplified by the finite size of the environmental partition. We also emphasise that the monotonicity of asymmetry measures follows from the assumption that the environment is initially uncorrelated with the system and prepared in a symmetry-respecting state (see Sec. E). In contrast, initialising the environment in a symmetry-breaking state, e.g., tilted ferromagnetic states, as done in Refs. [237, 238], violates this condition and can result in a non-G-covariant reduced map. This distinction accounts for the apparent non-monotonic behaviour observed in non-Markovian circuit realisations of the U(1)-symmetric Mpemba effect.

In the examples shown in Fig. 4.10(c-d), we observed that sectors with larger $|\mu|$ tend to decay faster. We expect that a similar phenomenon holds more broadly beyond the Markovian case discussed here, including for random U(1)-invariant circuits with an arbitrary number of layers. A better understanding of this phenomenon requires further investigation. Our preliminary numerics on such circuits also suggest connections between the quasiparticle picture (Refs. [232, 236]) and the mode-overlap picture,

since one can directly compare overlaps with slow quasiparticle carriers and with slow asymmetric eigenmodes.

As a final note, these non-Markovian circuits closely mirror setups in deep thermalisation with symmetries [261] and studies of dissipative quantum chaos [262].

Symmetry: SU(2)




Figure 4.12: Symmetry Mpemba effect for SU(2). The initial states are defined in Eq. (4.87), and we average over 100 random realisations of the dynamics generated by random SU(2)-symmetric gates with $N_s = 8$ and $N_e = 12$, drawing $J = J_z$ uniformly from $[-\pi/5, \pi/5]$ for each gate in $\hat{U}$.

We next turn to a non-abelian example, where the conserved quantity is the total angular momentum $\hat{\mathbf{S}}^2$. To impose full SU(2) covariance we switch off all local Z -rotations and Peierls phases ($h_n = h'_n = \phi_n = 0$) and sample only the isotropic exchange coupling $J = J_z$.


We break the symmetry on the system by preparing the state $\hat{\rho}_\theta = |\varphi(\theta)\rangle\langle\varphi(\theta)|$ with

$$|\varphi(\theta)\rangle = \cos(\theta/2) |\xi\rangle^{\otimes \lfloor N_s/2 \rfloor} + \sin(\theta/2) |0\rangle^{\otimes N_s}, \quad (4.87)$$

where $|\xi\rangle = (|01\rangle - |10\rangle)/\sqrt{2}$ is the singlet. For $\theta \neq k\pi \ \forall k \in \mathbb{Z}$, we find $|\varphi(\theta)\rangle$ in a superposition of the SU(2)-invariant singlet $|\xi\rangle^{\otimes \lfloor N_s/2 \rfloor}$ and the SU(2)-breaking product $|0\rangle^{\otimes N_s}$. The environment, which contains an even number of qubits N_e , is also initialised in singlets, as $|\xi\rangle^{\otimes N_e/2}$. Note that since the environment's state respects the SU(2) symmetry, and the Hamiltonian also respects this symmetry, the dynamics of the reduced system likewise respects the symmetry; that is, it is SU(2)-covariant. (When the environment contains an odd number of qubits, it must be prepared in a mixed state in order to obtain an SU(2)-invariant state.)

Fig. 4.12 shows five tilted states on an $N_s = 8$ system coupled to $N_e = 12$ environment qubits. As we vary θ , one can see characteristic crossings in their relaxation curves.

Figure 4.13: Splitting of the relative entropy of athermality for systems exhibiting an arbitrary global symmetry into the relative entropy of asymmetry and the relative entropy of athermality for symmetry-respecting states. An initially asymmetric state $\hat{\rho}$ evolves under a G -covariant map towards the steady state that lies within the set of G -invariant states (see Sec. 4.4.1). The relative entropy of athermality of $\hat{\rho}(t)$ ($\hat{\rho}_t$ in the figure, violet) can be split into two components: the relative entropy of $\hat{\rho}(t)$ with respect to the instantaneously dephased state $\mathcal{G}[\hat{\rho}(t)]$ (orange, namely the relative entropy of asymmetry) and the relative entropy of $\mathcal{G}[\hat{\rho}(t)]$ with respect to $\hat{\pi}$ (green, namely the relative entropy of athermality for symmetry-respecting states).



4.5 A unified description of the thermal and the symmetry Mpemba effects

We have reinterpreted instances of both the thermal (□) and symmetry (○) Mpemba effects in systems interacting with featureless baths as well as in non-Markovian environments using a resource-theoretic approach in both classical and quantum settings. Strikingly, the mechanisms behind the thermal and symmetry Mpemba effects are remarkably similar: the former is governed by the slowest-decaying mode, while the latter relies on the slowest symmetry-restoring mode. Importantly, in both cases, the relevant monotone is based on the quantum relative entropy; what differs is the reference state with respect to which it is computed. For the thermal Mpemba effect □, it characterises the divergence with respect to the steady state, while for the symmetry Mpemba effect ○, it quantifies the divergence with respect to an instantaneously dephased state. It is well-known that for the Davies map, the relative entropy between a non-equilibrium state and the thermal steady state can be split into a classical and a purely quantum coherent part [263]. In this context, the system's Hamiltonian is the symmetry generator and, correspondingly, the twirling operation consists of dephasing the non-equilibrium state in the energy eigenbasis. Here we show that this approach can be generalised to an arbitrary symmetry as follows: Let G be a group with unitary representation $\{\hat{U}_g\}_{g \in G}$ and $\mathcal{G}$ be the corresponding twirling map as in Eq. (4.42). For two states $\hat{\rho}, \hat{\pi}$ with $\mathcal{G}[\hat{\pi}] = \hat{\pi}$, the quantum relative entropy

reads

$$\begin{aligned}
 S(\hat{\rho} \parallel \hat{\pi}) &= \text{Tr}[\hat{\rho} \ln \hat{\rho}] - \text{Tr}[\hat{\rho} \ln \hat{\pi}] \\
 &\quad + \text{Tr}[\hat{\rho} \ln \mathcal{G}[\hat{\rho}]] - \text{Tr}[\hat{\rho} \ln \mathcal{G}[\hat{\rho}]] \\
 &= S(\hat{\rho} \parallel \mathcal{G}[\hat{\rho}]) - \text{Tr}[\hat{\rho} \ln \hat{\pi}] + \text{Tr}[\hat{\rho} \ln \mathcal{G}[\hat{\rho}]] .
 \end{aligned} \tag{4.88}$$

Now since $\hat{\pi}$ is G-invariant, we have that $\text{Tr}[\hat{\rho} \ln \hat{\pi}] = \text{Tr}[\hat{\rho} \ln \mathcal{G}[\hat{\pi}]]$. $\mathcal{G} \circ \mathcal{G} = \mathcal{G}$ and $\mathcal{G}^\dagger = \mathcal{G}$ imply

$$\text{Tr}[\hat{\rho} \ln \mathcal{G}[\hat{\pi}]] = \text{Tr}[\hat{\rho} \mathcal{G}[\ln(\mathcal{G}[\hat{\pi}])]] = \text{Tr}[\mathcal{G}[\hat{\rho}] \ln \hat{\pi}] . \tag{4.89}$$

Analogously, we have that $\text{Tr}[\hat{\rho} \ln \mathcal{G}[\hat{\rho}]] = \text{Tr}[\mathcal{G}[\hat{\rho}] \ln \mathcal{G}[\hat{\rho}]]$. Hence, we find that

$$S(\hat{\rho} \parallel \hat{\pi}) = S(\hat{\rho} \parallel \mathcal{G}[\hat{\rho}]) + S(\mathcal{G}[\hat{\rho}] \parallel \hat{\pi}) . \tag{4.90}$$


This equality holds for an arbitrary G-invariant state $\hat{\pi}$. In the special case where $\hat{\pi}$ is the thermal state, this identity implies that the relative entropy of athermality can always be separated into the relative entropy of asymmetry (symmetry-breaking contribution) and the relative entropy of athermality for symmetry-respecting states (symmetry-respecting contribution), as illustrated in Fig. 4.13. We note that, by Eq. (4.24), in the special case where $\hat{\pi} = \hat{\pi}_\beta$ is the thermal state at temperature β^{-1} , the relative entropy $S(\hat{\rho} \parallel \hat{\pi}_\beta)$ can be expressed as the difference between the non-equilibrium free energy of $\hat{\rho}$ and that of $\hat{\pi}_\beta$. Furthermore, by Eq. (4.49), which was shown in [255], we have $S(\mathcal{G}[\hat{\rho}] \parallel \hat{\rho}) = S(\mathcal{G}[\hat{\rho}]) - S(\hat{\rho})$. Hence, in this special case, Eq. (4.90) simplifies to

$$\begin{aligned}
 F(\hat{\rho}) - F(\hat{\pi}_\beta) &= \beta^{-1} (S(\mathcal{G}[\hat{\rho}]) - S(\hat{\rho})) \\
 &\quad + F(\mathcal{G}[\hat{\rho}]) - F(\hat{\pi}_\beta) ,
 \end{aligned} \tag{4.91}$$

which was previously presented in [256].

This decomposition reveals that for Davies maps, the quantum thermal Mpemba effect $\square$ comprises a purely classical component $\triangle$, corresponding to the classical thermal Mpemba effect, and a genuinely quantum part, which is associated with the restoration of time translational symmetry generated by $\hat{H}_s$ $\circ$. Here, as discussed in Ref. [194], eigenvalues associated with symmetry-preserving eigenmodes are real, while those corresponding to symmetry-breaking modes are complex (see Fig. 4.8). We stress that this distinction is not valid for more general G-covariant maps, as shown in Fig. 4.10. Moreover, there is nothing intrinsically quantum mechanical about the symmetry-breaking part of the dynamics. In fact, we showed a fully classical instance of a symmetry Mpemba effect in Sec. 4.4.4. So in general, Eq. (4.90) decomposes the total athermality (or non-stationarity) monotone into two parts: the asymmetry

Figure 4.14: Different Mpemba effects in the same thermalisation process described by the Davies map. As in Fig. 4.8, we study a TFIM, setting $J = 1$, $h = 1$, $N_s = 4$, $\beta = 2$, $\beta' = 1$ in both panels, and initialising the system in the state $\hat{\rho}_1 = \hat{\rho}_{\text{Gibbs}}^{\text{rand}} = (\hat{\rho}_{\text{Gibbs}} + \gamma \hat{\rho}_{\text{rand}})/N$, and $\hat{\rho}_2 = \hat{R}_{\text{metr}} \hat{\rho}_1 \hat{R}_{\text{metr}}^\dagger$ as in Sec. 4.4.6 (full lines). The dashed lines are obtained by stochastically minimising the overlap with the slowest-decaying mode via the Metropolis algorithm [194]. Panel (a): for weak noise $\gamma = 0.05$, the dynamics display a classical thermal $\triangle$ and a quantum thermal $\square$ Mpemba effect. Panel (b): for strong noise $\gamma = 0.25$ we find only the symmetry $\circ$ Mpemba effect.



monotone and a second part that quantifies the symmetry-respecting athermality⁹. As a final example, we study the occurrence of these different types of Mpemba effects in the same system. In Fig. 4.14, in both panels we consider the same thermalisation process described by a Davies map, but for different initialisations of the state $\hat{\rho}_1$ as the Gibbs state with random noise defined in Eq. (4.74) with $\gamma = 0.05$ (a) and $\gamma = 0.25$ (b), while as in Fig. 4.8, $\hat{\rho}_2 = \hat{R}_{\text{metr}} \hat{\rho}_1 \hat{R}_{\text{metr}}^\dagger$. Here, the left panel shows both the classical and quantum thermal Mpemba effects, while the right panel displays only the symmetry Mpemba effect. While a crossing in either $S(\hat{\rho}(t) \parallel \mathcal{G}[\hat{\rho}(t)])$ or $S(\mathcal{G}[\hat{\rho}(t)] \parallel \hat{\pi})$ is necessary to induce a crossing in the total relative entropy $S(\hat{\rho}(t) \parallel \hat{\pi})$, Eq. (4.90) implies that only when both component entropies cross we are guaranteed to find a crossing in $S(\hat{\rho}(t) \parallel \hat{\pi})$. This example demonstrates that in a given setting, one can seek different types of Mpemba effects, each connected to the depletion of a specific resource, by just changing the resource monotone M .

4.6 Complexity Mpemba effects

A notable recent variant of Mpemba phenomenology is the Pontus-Mpemba effect [201], which shows that, for a fixed initial state, an initial “heating” stage (driving the system further from equilibrium) can nonetheless outperform direct relaxation.

⁹The resource theory of symmetry-respecting athermality is defined by the free states $\mathcal{F} = \{\hat{\rho} \mid \mathcal{G}[\hat{\rho}] = \hat{\pi}\}$ and the free operations given by G -covariant TOs, such as the Davies map in our example. By construction, any free operation $\mathcal{E}$ maps free states to free states: G -covariance, $\mathcal{U}_g \circ \mathcal{E} = \mathcal{E} \circ \mathcal{U}_g$, implies $\mathcal{G} \circ \mathcal{E} = \mathcal{E} \circ \mathcal{G}$, so for $\hat{\rho} \in \mathcal{F}$ one has

$$\mathcal{G}[\mathcal{E}[\hat{\rho}]] = \mathcal{E}[\mathcal{G}[\hat{\rho}]] = \mathcal{E}[\hat{\pi}] = \hat{\pi}, \quad (4.92)$$

hence $\mathcal{E}[\hat{\rho}] \in \mathcal{F}$.




Figure 4.15: (a) Setup: a resourceful state $|\psi(\theta)\rangle$ is prepared, with θ controlling its resource content. For the Mpemba effect, different initial states evolve under free operations (grey); the resource monotone decreases locally in an effective system s (light orange) while the corresponding global resource monotone is conserved under global free unitaries. In the Pontus-Mpemba effect, the same state is first “heated” (by a resourceful stage) on s and e before free evolution. (b) Mpemba effects: A more resourceful state dissipates faster locally than a less resourceful one (crossing of monotones). (c) Pontus-Mpemba effects: An initial preheating stage accelerates local dissipation compared to free evolution alone.

Magic resources [264–267], non-Gaussianity [268–271], and imaginarity [272–275] are widely used as resource-theoretic quantifiers of non-classical computational power, and thus as proxies for quantum complexity. In the following sections we test the resource-theoretic Mpemba criterion for these complexity resources [4], which lack a direct thermodynamic or symmetry interpretation compared with the athermality and asymmetry measures considered above. We also show that, in these cases, under *global free* unitary dynamics, a more resourceful initial state can dissipate local resource faster than a less resourceful one, leading to crossings of a suitable subsystem monotone. We study this using the (subsystem) relative entropy of resource as monotone (or mana for the resource of magic), evaluated on reduced states of a closed system evolved under brickwork circuits of random two-body gates, starting from tilted product states. We find that imaginarity exhibits a Mpemba effect in this setting, whereas magic and non-Gaussianity do not (for the state families and free gate ensembles considered). Strikingly, a resourceful preheating stage induces a quantum Pontus-Mpemba effect for all resources considered here, with relaxation accelerated compared to direct free evolution from the same initial state.

To treat different resources on equal footing, we adopt a common circuit-based setup. Specifically, we consider a system of N qudits initialised in a resourceful product state, similar to Eq. (4.77), $|\psi_\alpha(\theta)\rangle = \bigotimes_{n=1}^N e^{-i\theta \hat{\sigma}_n^\alpha} |0\rangle_n$. Unlike before, we take an initial product state that is resourceful on any choice of s system, the global state carries nonzero resource uniformly while global free unitaries conserve it. This is the same initial-state family used, for example, in Refs. [199, 236, 237].

The dynamics is generated by a random brickwork circuit similar to the one defined in

Eq. (4.76). To tune the effective interaction density (and thereby control the Mpemba crossing time), the two-body gates are applied stochastically: with probability ϵ , $\hat{u}_{n,n+1}$ is drawn from the relevant subgroup of free operations for the resource under study, while with probability $1 - \epsilon$ the layer acts trivially as $\hat{I}_{n,n+1}$. Further, we consider either PBC or OBC. We consider a system $s = \{1, \dots, N_s\}$ and its complement e as an effective environment. Different setups are specified by three ingredients: (i) the family of initial states, (ii) the gate ensemble, and (iii) the resource monotone M . In this setting, the global resource content is conserved under global free unitaries, but the reduced states on s , $\hat{\rho}_\theta(t) = \text{Tr}_B[\hat{U}^t|\psi(\theta)\rangle\langle\psi(\theta)|\hat{U}^{-t}]$, follow a reduced dynamics that is not, in general, a free operation on A . When s is much smaller than e , the subsystem can effectively dissipate its local resource into e , which plays the role of an effective reservoir.

Beside the Mpemba effect, we define the Pontus-Mpemba effect [201], which arises in a complementary setting Fig. 4.15(d): starting from the same state $|\psi\rangle$, an additional resourceful “heating” stage $\hat{\tilde{U}}$ applied for T steps, followed by free evolution $\hat{U}^t$, can lead to faster resource dissipation than direct free evolution at the same total circuit depth $T + t$. Concretely, this occurs when (i) $M(\hat{\tilde{\rho}}(T)) > M(\hat{\rho}(T))$ after the heating stage, yet (ii) $M(\hat{\tilde{\rho}}(T + \tau)) < M(\hat{\rho}(T + \tau))$ at later times, for some $\tau > 0$, and

$$\hat{\tilde{\rho}}(t) = \begin{cases} \text{Tr}_e[\hat{\tilde{U}}^t|\psi\rangle\langle\psi|\hat{\tilde{U}}^{-t}] & t \leq T \\ \text{Tr}_e[\hat{U}^{t-T}\hat{\tilde{U}}^T|\psi\rangle\langle\psi|\hat{\tilde{U}}^{-T}\hat{U}^{T-t}] & t > T. \end{cases} \quad (4.93)$$

In the following sections, we demonstrate how these phenomena arise within the resource theories of imaginarity, non-Gaussianity, and magic. We start by introducing these three RTs in Secs. 4.6.1 to 4.6.3, then discuss the Mpemba effect in this context in Sec. 4.6.4 and conclude with the Pontus-Mpemba effect in Sec. 4.6.5.

4.6.1 Resource of imaginarity

Imaginarity quantifies the “complex phase” structure that is absent in the real (rebit) subtheory, where states and operations are represented by real matrices in a fixed basis. When control is restricted to real operations (orthogonal gates), access to imaginarity captures an operational overhead: generic complex quantum processes require additional ingredients (e.g. extra ancillas/controls), and imaginarity functions as a bookkeeping resource for this added implementation complexity [276–280].

Fix a reference basis $\{|i\rangle\}$ (typically the computational basis). The free states are those that are real in this basis, $\hat{\rho}$ free iff $\hat{\rho} = \hat{\rho}^*$ and free operations are real CPTP maps, i.e. channels admitting a Kraus representation with real matrices in the same

basis [276–278]. Introduce the (generally unphysical) “realification” map $\mathfrak{R}$ acting on matrix elements as $\langle i | \mathfrak{R}[\hat{\rho}] | j \rangle = \text{Re } \langle i | \hat{\rho} | j \rangle$. The relative entropy of imaginarity then takes the closed form [276]

$$M(\hat{\rho}) = S(\mathfrak{R}[\hat{\rho}]) - S(\hat{\rho}), \quad (4.94)$$

which is additive under tensor products, and it is monotone under real operations [276–278]. In our numerics the global evolution is generated by two-body free (orthogonal) gates

$$\hat{u} = \exp \left[-2i\alpha (\hat{s}^x \otimes \hat{s}^y - \hat{s}^y \otimes \hat{s}^x) \right] \in \text{SO}(4) \subset \text{U}(4), \quad (4.95)$$

with α sampled uniformly from $[0, 2\pi]$. Since the generator is purely imaginary in the computational basis, $\hat{u}$ is a real unitary (orthogonal) gate and therefore cannot create imaginarity from globally real initial conditions.

4.6.2 Resource of non-Gaussianity

Fermionic Gaussian states and Gaussian (free-fermion) evolutions form a classically tractable corner of many-body quantum dynamics: they admit an efficient description in terms of covariance matrices and underpin the classical simulability of free-fermion (matchgate) circuits [281, 282]. Departures from Gaussianity capture the additional structure generated by interactions or measurements, and are therefore a natural *complexity resource* for going beyond free-fermion simulation; in this spirit, pure fermionic non-Gaussianity is closely related to “magic” for matchgate computations [269, 283].

The free states are fermionic Gaussian states (for parity-preserving states, fully specified by their two-point correlators), and the free unitaries are fermionic Gaussian unitaries generated by quadratic Hamiltonians (interactions or measurements can create non-Gaussian correlations and thus increase the resource). On a spin-1/2 chain it is convenient to introduce $2N$ Majorana operators via the Jordan-Wigner map

$$\hat{\gamma}_{2a-1} = \left(\prod_{i=1}^{a-1} \hat{\sigma}_i^z \right) \hat{\sigma}_a^x, \quad \hat{\gamma}_{2a} = \left(\prod_{i=1}^{a-1} \hat{\sigma}_i^z \right) \hat{\sigma}_a^y, \quad (4.96)$$

which satisfy $\{\hat{\gamma}_a, \hat{\gamma}_b\} = 2\delta_{a,b}$. For a reduced state $\hat{\rho}_s$ on N_s sites (hence $2N_s$ Majoranas), define the antisymmetric covariance matrix

$$\Gamma_{a,b}(\hat{\rho}_s) = -\frac{i}{2} \text{Tr}(\hat{\rho}_s [\hat{\gamma}_a, \hat{\gamma}_b]). \quad (4.97)$$

Let $\hat{\rho}_s^{\mathcal{G}}$ be the (unique) fermionic Gaussian state with the same covariance matrix $\Gamma(\hat{\rho}_s)$; equivalently, it is the maximum-entropy state compatible with these second moments. In this case, Eq. (4.7) becomes the relative entropy distance to the Gaussian

set (relative entropy of non-Gaussianity) [268, 284]

$$M(\hat{\rho}_s) = S(\hat{\rho}_s^{\mathcal{G}}) - S(\hat{\rho}_s). \quad (4.98)$$

In practice, $S(\hat{\rho}_s^{\mathcal{G}})$ follows directly from $\Gamma(\hat{\rho}_s)$: if the spectrum of Γ is $\{\pm\lambda_j\}_{j=1}^{N_s}$ with $\lambda_j \in [0, 1]$, then

$$S(\hat{\rho}_s^{\mathcal{G}}) = \sum_{j=1}^{N_s} H\left(\frac{1+\lambda_j}{2}\right), \quad H(x) = -x \log_2 x - (1-x) \log_2 (1-x). \quad (4.99)$$

In our numerics the free evolution is built from two-site Gaussian gates acting on four Majoranas, i.e. unitaries generated by quadratic (free-fermion) forms

$$\hat{u} = \exp\left[\frac{1}{4} \sum_{a,b=1}^4 K_{a,b} \hat{\gamma}_a \hat{\gamma}_b\right], \quad K = -K^T \text{ real}. \quad (4.100)$$

4.6.3 Resource of magic

The resource theory of magic formalises the gap between the efficiently classically simulable stabiliser subtheory and universal quantum computation. Let $\mathcal{P}_N$ be the N -qubit Pauli group (tensor products of $\{\hat{I}, \hat{\sigma}^x, \hat{\sigma}^y, \hat{\sigma}^z\}$ with phases), and $\mathcal{C}_N$ the Clifford group (the normaliser of $\mathcal{P}_N$),

$$\mathcal{C}_N = \left\{ \hat{U} \in \text{U}(2^N) : \hat{U} \hat{P} \hat{U}^\dagger \in \mathcal{P}_N \quad \forall \hat{P} \in \mathcal{P}_N \right\}. \quad (4.101)$$

Pure stabiliser states are generated from computational-basis product states by Clifford circuits,

$$|\phi\rangle \in \text{STAB}_{\text{pure}} \iff \exists \hat{U} \in \mathcal{C}_N : |\phi\rangle = \hat{U} |0\rangle^{\otimes N}, \quad (4.102)$$

and the free (mixed) states form the stabiliser polytope,

$$\text{STAB} = \text{conv}(\{|\phi\rangle\langle\phi| : |\phi\rangle \in \text{STAB}_{\text{pure}}\}). \quad (4.103)$$

The Clifford/stabiliser sector admits efficient classical simulation (Gottesman-Knill and refinements), so STAB is treated as free [285]. Non-stabiliser (“magical”) states enable universality via state injection [264]. Mixed-state magic is subtle: some $\hat{\rho} \notin \text{STAB}$ are not distillable (“bound magic”) [286].

A distance-based monotone is the relative entropy of magic,

$$M_{\text{rel}}(\hat{\rho}) = \min_{\hat{\sigma} \in \text{STAB}} S(\hat{\rho} \parallel \hat{\sigma}), \quad S(\hat{\rho} \parallel \hat{\sigma}) = \text{Tr}[\hat{\rho}(\log \hat{\rho} - \log \hat{\sigma})], \quad (4.104)$$

but the minimisation over STAB is typically intractable at many-body sizes.

For odd prime local dimension d (in particular $d = 3$ qutrits), a widely used and directly computable alternative is the *mana*, defined from the Gross discrete Wigner representation. Let $\{A_{\mathbf{r}}\}_{\mathbf{r} \in \mathbb{Z}_d^{2N}}$ be the phase-point operators and

$$\mathcal{W}_{\mathbf{r}}[\hat{\rho}] = \frac{1}{d^N} \text{Tr}[\hat{\rho} \hat{A}_{\mathbf{r}}] \quad \left(\sum_{\mathbf{r}} \mathcal{W}_{\mathbf{r}} = 1 \right). \quad (4.105)$$

The mana is the log of the ℓ_1 -norm of $\mathcal{W}$,

$$M(\hat{\rho}) = \log_2 \sum_{\mathbf{r} \in \mathbb{Z}_d^{2N}} |\mathcal{W}_{\mathbf{r}}[\hat{\rho}]|, \quad (4.106)$$

which is additive under tensor products and monotone under stabiliser operations; for odd d , Wigner-positivity is tightly linked to stabiliserness [286–288].

In this thesis we focus on a one-dimensional chain of N qutrits ($d = 3$). With generalised Pauli operators

$$\hat{Z} = \sum_{m=0}^{d-1} e^{2\pi i m/d} |m\rangle\langle m|, \quad \hat{X} = \sum_{m=0}^{d-1} |m\rangle\langle m+1 \pmod{d}|, \quad (4.107)$$

we write Pauli strings as $\hat{P}_{\mathbf{r}} = \prod_{l=1}^N e^{\pi i (d+1) r_l^x r_l^z / 2} \hat{X}_l^{r_l^x} \hat{Z}_l^{r_l^z}$ with $\mathbf{r} = (r_1^x, r_1^z, \dots, r_N^x, r_N^z) \in \mathbb{Z}_d^{2N}$ and define $\hat{A}_0 = \sum_{\mathbf{r}} \hat{P}_{\mathbf{r}} / d^N$, $\hat{A}_{\mathbf{r}} = \hat{P}_{\mathbf{r}} \hat{A}_0 \hat{P}_{\mathbf{r}}^\dagger$ [287]. Starting from the tilted ferromagnetic state

$$|\psi_{\tilde{x}}(\theta)\rangle = \bigotimes_{n=1}^N e^{-i\theta(\hat{\sigma}_n^x + \hat{\sigma}_n^{x\dagger})} |0\rangle_n, \quad (4.108)$$

where the $\tilde{x}$ distinguishes it from the state $|\psi_x(\theta)\rangle$, we study the evolution of the subsystem mana $M(\hat{\rho}_s)$ for N_s contiguous qutrits under a global brickwork circuit composed of two-qutrit Clifford gates $\hat{u} \in \mathcal{C}_2(d)$.

4.6.4 Mpemba effects of complexity

Keeping in mind Sec. 4.6.1, we begin with the resource theory of imaginarity [272–275, 289].

The numerical results are shown in Fig. 4.16(a) for a chain of $N = 20$ qubits and a system of size $N_s = 2$, averaged over 4000 circuit realisations with $\epsilon = 1$, starting from the x -tilted product state $|\psi_x(\theta)\rangle = \bigotimes_{n=1}^N e^{-i\theta \hat{\sigma}_n^x} |0\rangle_n$.

The relative entropy of imaginarity on s exhibits the hallmark crossing behaviour: larger θ corresponds to a more resourceful initial reduced state, yet the corresponding $M(\hat{\rho}_\theta(t))$ decays faster with circuit depth, yielding a clear Mpemba crossing .

The other resources analysed do not exhibit Mpemba crossings for these tilted product-state families. We illustrate this with two examples: (i) non-Gaussianity [268–271]




Figure 4.16: Mpemba diagnostics for complexity resources in 1D circuits: imaginarity (a) and non-Gaussianity (b) for $N=20$ qubits, and magic (mana) (c) for $N=12$ qutrits. We evaluate the resource content of the system s of size $N_s=2$ using the appropriate monotones, starting from tilted product states $|\psi_\alpha(\theta)\rangle$. The dynamics is generated by a depth- $2t$ circuit $\hat{U}^t$ composed of resource-free two-body gates (applied with dilution ϵ), and results are averaged over 4000 circuit realisations. For the chosen state families, clear crossings occur for imaginarity (a) $\heartsuit$, while no comparable crossings are observed for non-Gaussianity or magic (b–c).

and (ii) magic, quantified via mana [286].

For non-Gaussianity (Fig. 4.16(b)) we repeat the same protocol as for imaginarity, same tilted family and the corresponding free gate ensemble, yet observe no crossings of $M(\hat{\rho}_\theta(t))$ within numerical resolution for the explored θ . As in the previous cases, the global mana is conserved under global free operations, while the reduced state on the system s can lose resource as the circuit redistributes it into the effective environment e . Our numerical results for dilution $\epsilon = 0.5$, with $N = 12$ qutrits and $N_s = 2$, are shown in Fig. 4.16(c). As for non-Gaussianity, the mana $M(\hat{\rho}_\theta(t))$ shows no crossings for the explored tilts, indicating that these tilted product states do not display a Mpemba effect for magic in this setting.

In summary, imaginarity exhibits clear Mpemba crossings for the tilted families considered here, whereas non-Gaussianity and magic do not. Remarkably, even in these latter cases a resourceful preheating stage can induce crossings and accelerate local relaxation, realising a Pontus-Mpemba effect, as discussed next.

4.6.5 Pontus-Mpemba Effects

We now fix a resourceful initial state $|\psi(\theta)\rangle$ and study how a preheating stage modifies the subsequent free relaxation of the system s . For circuit depth $t \leq T$, we first apply a resourceful “heating” circuit that acts independently on the system s and effective environment e ,

$$|\tilde{\psi}(t)\rangle = (\hat{\tilde{U}}_s^t \otimes \hat{\tilde{U}}_e^t) |\psi(\theta)\rangle, \quad (4.109)$$




Figure 4.17: Pontus-Mpemba diagnostics for complexity resources. (a) Imaginarity and (b) Non-Gaussianity for $N = 20$ qubits ($d = 2$) and (c) magic (mana) for $N = 12$ qutrits ($d = 3$), all with system size $N_s = 2$. Starting from $|\psi(\theta)\rangle$, we apply a two-step protocol: (i) a resourceful preheating circuit $\hat{U}_s^T \otimes \hat{U}_e^T$ of depth $2T$, followed by (ii) a free circuit $\hat{U}^{t-T}$. All results are averaged over 4000 circuit realisations. Increasing T leads to faster decay of the local resource content on s , consistent with the Pontus-Mpemba effect.

where $\hat{U}_{s/e}$ are depth-2 brickwork circuits composed of two-body gates $\hat{u}$ sampled independently from the Haar measure on $U(d^2)$ at each time step t . For convenience, we introduce two dilution parameters, ϵ_s and ϵ_e , controlling the density of Haar gates in the preheating on s and e , respectively, applying a Haar gate with probability $\epsilon_{s/e}$, else the identity. After the heating stage, the dynamics is generated by the free circuit $\hat{U}^{t-T}$, i.e. the same resource-free gate ensemble as in Fig. 4.15(a).

We first consider imaginarity. We fix the initial tilt to $\theta = 0.55$ and apply a two-step protocol (preheating on s and e , then free evolution). In Fig. 4.17(a) we set $\epsilon_s = 0.55$ and the free-circuit dilution to $\epsilon = 0.6$, while scanning the preheating depth $0 \leq T \leq 4$ and the environment preheating strength $0 \leq \epsilon_e \leq 0.6$. As expected, preheating increases the local resource content on the system s up to $t = T$; once heating is switched off, the subsequent dynamics is purely free. Once free evolution begins, the prior preheating accelerates the decay of non-Gaussianity $M_s(t)$: clear crossings appear (♥) as T and (ϵ_s, ϵ_e) are varied, realising Pontus-Mpemba phenomenology [201] in a unitary circuit setting.

We next turn to non-Gaussianity. We fix $|\psi(\theta)\rangle$ at $\theta = 0.7$ and repeat the same two-step protocol: preheating for $t \leq T$ followed by free evolution for $t > T$. The results in Fig. 4.17(b) use $\epsilon_s = 0.25$ and $\epsilon = 0.9$, and scan $0 \leq T \leq 3$ and $0 \leq \epsilon_e \leq 1$ for $N = 20$ qubits ($d = 2$) with $N_s = 2$. Again, preheating increases the local resource content on s up to $t = T$, but once it is turned off the heated state locally dissipates the resource faster (♦).

Finally, we consider magic for $N = 12$ qutrits with $N_s = 2$. Fixing $|\psi_{\vec{x}}(\theta)\rangle$, we compare $T = 0, 1$ at $\epsilon_s = 0.01$, $\epsilon_e = 1.0$, and $\epsilon = 0.4$, see Fig. 4.17(c). Once more, the two-step (resourceful + free) protocol yields faster relaxation of $\hat{\rho}_s$, and hence of the monotone

M_s , than purely free evolution, consistent with the Pontus-Mpemba effect (☆). Taken together, these results indicate that the Pontus-Mpemba effect is broadly present across the complexity resource theories studied here: when ordinary Mpemba crossings are absent for a given state family, preheating can induce them, and when they are present, preheating preserves, and can even enhance them.

4.7 Conclusion

This chapter framed Mpemba physics as anomalously fast depletion of a resource under free dynamics: a Mpemba effect occurs when the monotone of a more resourceful initial state drops below that of a less resourceful one, producing a crossing. Within this viewpoint, thermal Mpemba effects correspond to the decay of athermality (non-equilibrium free energy), while symmetry Mpemba effects correspond to the decay of asymmetry, which we analysed via modes of asymmetry and their overlaps with slow relaxation modes. The unified relative-entropy decomposition in [Sec. 4.5](#) further clarifies how different monotones isolate different relaxation channels, so that distinct Mpemba effects can coexist within the same physical dynamics. From a resource-theoretic perspective, a general criterion for the occurrence of a Mpemba effect, with respect to a given resource and its associated free operations, is the existence of a pair of states such that the more resourceful one has a smaller overlap with the slowest resourceful relaxation eigenmode than the less resourceful one. Using this perspective, we identified a simple classical model exhibiting a Mpemba effect, thereby demonstrating that a symmetry Mpemba effect can arise in a purely classical setting.

Finally, we extended the same resource-theoretic criterion to complexity resources in closed many-body circuits, where global resource conservation under free unitaries coexists with local dissipation on subsystems. In this setting, Pontus-Mpemba protocols provide a systematic speedup mechanism: a resourceful preheating stage can accelerate the subsequent local relaxation even when the ordinary Mpemba crossings are absent for the same state family.

Part II

Characterising complexity

Complexity, chaos, and transport




Figure II.1: Different complexity types: **Part I** probes (i)-(v) mapped to spectral, transport, and operator complexity, and to the operator-growth diagnostics used in **Part II** (OTOCs, operator entanglement, and Krylov complexity). Dashed arrows indicate indirect/conditional links.

In many-body quantum physics, extracting emergent phenomena, which observables relax quickly, which stay frozen, how initial structure redistributes across space and time, typically requires identifying the right effective description. The central organising question is therefore how *structure* builds up and decays under evolution, and which reduced descriptions remain predictive at which scales. *Complexity* is not meant to be a single quantity, but rather a family of diagnostics that quantify how hard it is to describe, predict, or compress an evolving state, channel, or operator relative to a chosen reference class.

The role of this opening section is therefore best understood as a retrospective dictionary. Having established in **Part I** a set of concrete probes of spectral structure, conservation-law constraints, and real-time spreading in spin chains, Floquet circuits, and resource-theoretic settings, we can now translate those results into a complexity language. In that translation, “chaos” diagnostics fix the late-time ceiling, transport diagnostics identify the bottlenecks and protected sectors that shape the approach to it, and operator-growth diagnostics quantify how quickly simple operators become hard to represent. Making this dictionary explicit here also clarifies what information from **Part I** will be taken as *input* for the operator-space constructions developed next.

From a complexity viewpoint, the diagnostics of **Part I** naturally separate into three different notions: *spectral complexity* (how scrambled/incompressible is the spectrum/eigenbasis at late time), *transport complexity* (what slow modes/conservation laws bottleneck relaxation and operator growth), and *operator complexity* (how fast a simple operator becomes hard to describe in operator space), as depicted in **Fig. II.1**. Spectral universality sets the late-time ceiling (“saturated complexity”), while conservation laws and slow modes can strongly constrain the route to that ceiling, even in

spectrally chaotic dynamics. Fig. II.1 summarises how the Part I probes (i)-(v) feed into these notions and into the operator-growth diagnostics emphasised in Part II.

In Part II we focus on operator complexity. This is usually probed through three main diagnostics: OTOCs, local-operator entanglement (LOE), and *Krylov complexity*. OTOCs probe operator spreading through commutator growth [48, 290], LOE measures the entanglement of an operator $\hat{O}(t)$ in the Choi representation [291, 292], and Krylov complexity tracks how $\hat{O}(t)$ spreads across a Liouvillian-generated orthonormal operator basis [293–295]. Under the operator growth hypothesis (OGH), linear Lanczos growth $b_n \sim \alpha n$ bounds the OTOC Lyapunov exponent $\lambda_L \leq 2\alpha$ [293], showing how a chaotic growth mechanism in operator space (linear Lanczos growth) constrains the exponential growth of commutators, but not vice versa. Krylov complexity is explicitly seed-resolved and connects directly to moments and continued-fraction representations of correlators, linking back to the DOS/LDOS tools developed in Ch. 2.

Part I, instead, introduced spectral and dynamical probes for spin chains, Floquet circuits and resource-theoretic frameworks (Chs. 2 to 4), with emphasis on finite- and infinite-time behaviour. Spectral and eigenstate universality serve as diagnostics of *chaos*, whereas real-time correlators and spreading exponents diagnose *transport* and the hydrodynamic constraints that control it. Resource monotones quantify distance to a chosen free set and thus track irreversible loss of structure. The point of the present section is to make the dependency explicit: those probes already constrain *which* notion of complexity is relevant in a given regime, and they determine what one should expect for operator-growth diagnostics before introducing any Krylov machinery. Accordingly, we first distil the lessons of Part I into this complexity dictionary, and then turn to the technical construction in operator space.

(i) Eigenphase statistics, analysed in Sec. 2.1.1, are primarily a *chaos* diagnostic: the distinction between Wigner-Dyson and Poisson statistics distinguishes ergodic from integrable/localised regimes and signals whether the system has “forgotten” microscopic structure in its spectral correlations [45, 46]. From a complexity standpoint, RMT statistics indicate *saturated* spectral complexity at infinite time (maximally scrambled, incompressible spectra up to symmetry constraints), but level statistics alone do not determine *how fast* that ceiling is approached, nor which slow modes control intermediate-time growth.

(ii) Operator-resolved spectral measures (DOS, LDOS), discussed in Sec. 2.2, bridge *chaos* and *complexity*. The total DOS $g(E)$ sets the available spectral phase space, while the LDOS $g_\psi(E)$ reveals how a simple initial state decomposes in the energy eigenbasis [46]. Smooth, featureless LDOS is compatible with ETH-type thermalisation [59, 73] (see Sec. 2.1.2), while sharp features signal constraints or integrability;

Eigenphase
statistics

Spectral
measures

either way, these objects primarily diagnose state/eigenbasis structure rather than operator growth.

Current
spectral
functions

(iii) Current spectral functions $S_{JJ}(\omega)$, see [Sec. 2.3](#), are a *transport* diagnostic, with indirect complexity information. From a complexity viewpoint, a Drude contribution indicates operator components protected by conservation laws that do *not* decay (Mazur bound), remaining “simple” indefinitely, while the regular low-frequency continuum reflects components that do spread and complexify [\[43, 296\]](#). The balance between Drude weight and regular conductivity therefore exposes a hierarchy of slow/simple versus fast/complexifying sectors.

Transport
exponents

(iv) Transport exponents quantify the spatial spreading of conserved charge. The scaling $m_2(t) \sim t^\alpha$ classifies transport as ballistic ($\alpha = 2$), diffusive ($\alpha = 1$), subdiffusive ($0 < \alpha < 1$), or superdiffusive ($1 < \alpha < 2$) [\[43\]](#). They constrain how fast initially localised structure spreads in real space, but they are not by themselves a measure of operator-space complexity. Anomalous exponents and long prethermal regimes (as observed in the “swappy” regime of [Ch. 3](#)) instead serve as useful signatures for intermediate-time structure that can strongly shape operator-growth diagnostics [\[35, 109\]](#).

Resource-
monotone
decay

(v) Resource-monotone decay, discussed in [Ch. 4](#), is a *complexity*-theoretic reframing of equilibration. Relative-entropy distance to the free set, $\min_{\hat{\sigma} \in \mathcal{F}} S(\hat{\rho} \parallel \hat{\sigma})$, quantifies how far a state is from a “free” (structureless) reference [\[3, 4, 203, 255\]](#). Decay of these monotones under free operations directly parallels complexity saturation: initial resourcefulness (= structure) decays towards the free-state set (= simplicity). Modes of asymmetry [\[204\]](#) provide a spectral decomposition of this decay, identifying which symmetry-breaking components relax slowly versus quickly, a hierarchy directly analogous to the slow-mode/fast-mode distinction in transport. The Mpemba effect (faster relaxation from more-resourceful initial states) arises when strongly asymmetric states have small overlap with the slowest mode [\[174, 188, 199\]](#), paralleling how certain initial operators in Krylov constructions can bypass slow-growing subspaces and equilibrate faster despite higher initial complexity.

Krylov
complexity

Krylov complexity reduces operator growth to motion on a one-dimensional Krylov chain generated from a choice of inner product and a seed operator. Its explicit seed dependence separates universal growth laws from operator-specific transients, and its continued-fraction/moment structure links operator dynamics to equilibrium correlators.

[Part II](#) develops this operator-complexity framework explicitly, with an emphasis on Krylov complexity and the additional subtleties raised in Floquet settings: quasienergy

periodicity, branch (unfolding) choices, and the need for an effective generator when comparing continuous-time and discrete-time complexity data. In the next chapter we introduce the Krylov construction in operator space and develop the tools needed to apply it consistently to Floquet circuits.



Chapter 5

Krylov Complexity

5.1 Introduction

Krylov complexity compresses dynamics by projecting onto the Krylov subspace generated by a “seed” (state or operator) under repeated application of a linear generator (Hamiltonian/Liouvillian or Floquet channel). This subspace encodes seed-resolved moments of the dynamics, equivalently the moments of the associated spectral measure with respect to a chosen inner product; throughout we use the Hilbert-Schmidt product (infinite temperature). For Hermitian generators, Lanczos tridiagonalisation produces an orthonormal Krylov basis in which the generator becomes a one-dimensional nearest-neighbour problem. The resulting recursion (Lanczos) coefficients therefore act as a compact set of dynamical data: they determine seed-resolved spectral quantities (e.g. the LDOS and diagonal Green’s functions [297, 298], Sec. 2.2) and time-domain correlators [299–301].

This representation also gives an immediate geometric picture. In the Krylov basis the evolved seed is a wavepacket over the Krylov index, and the Krylov complexity $K(t)$ is (roughly) its mean position; growth of $K(t)$ quantifies the propagation into increasingly “complicated” Krylov directions.

Because practical Krylov methods truncate the basis at finite dimension, the same wavepacket picture provides a notion of simulability: slow growth of $K(t)$ suggests that $\hat{O}(t)$ remains confined to a low-dimensional region of operator space and can be efficiently approximated, whereas rapid growth indicates that larger Krylov subspaces are needed and truncations break down earlier.

In modern quantum chaos the emphasis has shifted from spectra to operator dynamics: scrambling probes such as OTOCs and “operator hydrodynamics” track the spread of initially local operators and the constraints imposed by conservation laws [119]. In the same spirit, Krylov complexity fits this operator-space viewpoint while retaining the Lanczos structure: iterating the dynamical map on an operator $\hat{O}$ yields a tridiagonal reduced dynamics specified by recursion data that coincide with the continued-fraction coefficients of equilibrium correlators (Mori formalism) [293–295, 300, 302]. This links Krylov recursion to moments of operator autocorrelators and, in principle, to time-domain measurements, subject to finite-time and noise limitations [295, 303].

Under the OGH, the Lanczos coefficients grow as $b_n \sim \alpha n$, which upper-bounds the Lyapunov exponent as $\lambda_L \leq 2\alpha$; we review the recursion mechanism in [Sec. 5.2.1](#). For spatially local, bounded Hamiltonians (or Liouvillians), one can bound the growth of b_n to be at most linear in n (with model-dependent constants) [\[293\]](#), so the OGH may be viewed as the statement that generic nonintegrable (thermalising) dynamics *saturates* this locality-limited growth (up to possible logarithmic corrections in 1D), while integrable/free models can display slower or more structured behaviour. This perspective will be useful below when discussing Floquet generators: if a chosen branch of $\log \hat{U}$ produces a highly nonlocal effective $\hat{H}_F$, the assumptions behind such locality bounds (and hence the OGH intuition) are no longer operative. Krylov complexity is best viewed as a *probe* rather than a definition of chaos: it is seed-resolved, its saturation value is not universal across integrable-chaotic crossovers [\[304\]](#), clear exponential regimes are not generic outside semiclassical limits [\[305\]](#), and in nonergodic phases such as MBL the effective Krylov chain can exhibit localisation, with the Lanczos coefficients growing linearly with a staggered second order correction [\[306\]](#).

A key advantage of the Krylov construction is that it is formulated directly in terms of the dynamical map on operator space. This makes the extension to stroboscopic Floquet circuits algebraically straightforward: one replaces the Liouvillian by the adjoint Floquet channel $\mathcal{U}[\cdot] = \hat{U}(\cdot)\hat{U}^\dagger$ and studies Krylov growth versus integer cycle number. In practice the reduced dynamics is typically constructed via Arnoldi iteration (the analogue of Lanczos for the non-Hermitian superoperator $\mathcal{U}$) [\[307–309\]](#). However, quasienergy periodicity and the associated many-to-one spectral folding mean that Arnoldi data can be less sensitive to microscopic structure than Lanczos data: for dual-unitary circuit families, [Ref. \[309\]](#) reports that $K_{\text{Arn}}(t)$ exhibits essentially the same maximal linear growth across noninteracting, integrable, and chaotic gates, in contrast with the OGH mechanism based on unbounded Lanczos growth. We therefore treat Krylov complexity as an operator-level diagnostic to be benchmarked against spectral (RMT) and operator-spreading probes within the same circuits.

The remainder of the Chapter is organised as follows. In [Sec. 5.2](#) we place Lanczos (continuous time) and Arnoldi (discrete time) constructions on the same footing and isolate the kinematic constraint that unitarity imposes on stroboscopic Krylov growth. In [Sec. 5.3](#) we then develop the spectral-measure bridge between continuous and discrete time, deriving the exact aliasing (folding) relation and illustrating its consequences in two controlled benchmarks ([Secs. 5.3.1](#) and [5.3.2](#)). The next block, [Secs. 5.4](#) and [5.5](#), addresses the central obstruction revealed by aliasing, the branch ambiguity of $\log \hat{U}$, and introduces a locality-guided adiabatic unfolding criterion (MEO) for selecting physically meaningful Floquet generators along a path in unitary space. We benchmark this strategy analytically and numerically in the brickwork-driven XX chain in [Sec. 5.6](#) and then apply the same pipeline to other dual-unitary endpoint families in [Sec. 5.7](#),

before summarising the implications for Krylov diagnostics in [Sec. 5.8](#).

The key takeaways, which will be made precise in the following sections, are the following. First, spectral aliasing is the fundamental reason there is no canonical one-to-one correspondence between Lanczos and Arnoldi recursion data extracted from stroboscopic sampling. Second, under a conservative principal-branch reconstruction, a maximally ergodic unitary (MEU)-type Floquet fixed point corresponds to a band-limited continuous-time measure with bounded Lanczos coefficients, and therefore does not support sustained OGH-type exponential $K_{\text{Lan}}(t)$. Third, when a meaningful unfolded Floquet generator is selected (here via MEO), one can recover smooth quasienergy continuations and substantially improve the locality of the effective Hamiltonian compared to the principal branch, enabling a controlled Lanczos and Krylov analysis of Floquet circuits. Finally, at dual-unitary endpoints the resulting Krylov behaviour is model-dependent: in free-fermion settings (XX) complexity is bounded by finite-dimensional Krylov orbits, while interacting families can display qualitatively different trends, as discussed in [Secs. 5.7](#) and [5.8](#).

Within this programme, Ref. [\[6\]](#) (in preparation) contains the original contributions developed during this PhD. Concretely, it (i) formulates a unified spectral-aliasing framework that connects continuous-time (Lanczos) and discrete-time (Arnoldi) Krylov data at the level of seed-resolved spectral measures and moments; (ii) introduces the locality-guided adiabatic unfolding protocol (MEO) as a practical branch-selection principle for $\log \hat{U}$ along a Floquet path; and (iii) applies the resulting unfolded generators to a systematic Krylov-complexity study across integrable and chaotic dual-unitary endpoint circuit families, establishing the structural split between bounded free-fermion dynamics and interacting behaviour highlighted in this Chapter.

5.2 Krylov bases: continuous vs. discrete time

Krylov constructions turn operator dynamics into single-particle motion on a semi-infinite “Krylov chain”: starting from a simple seed $\hat{O}_0$, repeated action of the dynamical map generates a nested Krylov subspace whose orthonormal basis $\{\hat{O}_m\}_{m \geq 0}$ resolves both correlation functions and a natural notion of operator spreading in Krylov space (“Krylov complexity”). For Hamiltonian dynamics the relevant generator is (a Hermitian) Liouvillian, and the associated orthonormalisation is the Lanczos tridiagonalisation. For Floquet circuit dynamics the natural object is instead the one-step unitary superoperator, and the corresponding orthonormalisation is an Arnoldi-type construction. We keep a unified notation while emphasising the different algebraic constraints (Hermiticity vs. unitarity).

	Lanczos (continuous time)	Arnoldi (discrete time)
Dynamics	$ O(t)\rangle = e^{i\mathcal{L}t} O_0\rangle$	$ O(t)\rangle = \mathcal{U}^t O_0\rangle, t \in \mathbb{Z}$
Map on operators	$\mathcal{L} = [\hat{H}, \cdot]$	$\mathcal{U}[\hat{O}] = \hat{U}^\dagger \hat{O} \hat{U}$
Spectrum	real line	unit circle
Krylov basis	Gram-Schmidt on $\{\mathcal{L}^m O_0\rangle\}$	Gram-Schmidt on $\{\mathcal{U}^m O_0\rangle\}$
Matrix form in Krylov basis	symmetric tridiagonal	upper-Hessenberg
Recursion coefficients	$b_m^{(\mathcal{L})} \geq 0$	$a_m^{(\mathcal{U})}, b_m^{(\mathcal{U})}, c_m^{(\mathcal{U})} \in [0, 1]$
Krylov complexity	$K_{\text{Lan}}(t) = \sum_m m \phi_m^{(\mathcal{L})}(t) ^2$	$K_{\text{Arn}}(t) = \sum_m m \phi_m^{(\mathcal{U})}(t) ^2 \leq t$
Generic late-time scaling	OGH: $b_m^{(\mathcal{L})} \sim \alpha m \Rightarrow K_{\text{Lan}}(t) \sim e^{2\alpha t}$	MEU regime: $b_m^{(\mathcal{U})} \rightarrow 1, a_m^{(\mathcal{U})} \rightarrow 0$ (and left-hopping terms vanish) $\Rightarrow K_{\text{Arn}}(t) \sim t$

Table 5.1: Parallel summary of Lanczos (Hamiltonian/Liouvillian) and Arnoldi-type Krylov constructions (Floquet/unitary superoperator), highlighting the distinct constraints on recursion data and the resulting asymptotic growth of Krylov complexity.

We regard operators as vectors equipped with the inner product

$$(A|B) = \frac{1}{D} \text{Tr}[\hat{A}^\dagger \hat{B}]. \quad (5.1)$$

The associated norm is

$$\|A\| = \sqrt{(A|A)} = \sqrt{\frac{1}{D} \text{Tr}[\hat{A}^\dagger \hat{A}]} . \quad (5.2)$$

Choosing a normalised seed $|O_0\rangle$, we construct an orthonormal Krylov basis $\{|O_m\rangle\}_{m \geq 0}$ and expand the evolved operator as

$$|O(t)\rangle = \sum_{m \geq 0} \phi_m(t) |O_m\rangle, \quad \phi_m(t) = (O_m|O(t)). \quad (5.3)$$

Throughout, $|\cdot\rangle$ denotes operators viewed as vectors in operator space, with duals $(\cdot|$. Continuous-time dynamics uses $t \in \mathbb{R}$, while discrete-time (Floquet) dynamics uses integer steps $t \in \mathbb{Z}$. Normalisation of $|O(t)\rangle$ implies $\sum_{m \geq 0} |\phi_m(t)|^2 = 1$, so $|\phi_m(t)|^2$ defines a probability distribution over the Krylov index m . The Krylov complexity is the first moment of this distribution,

$$K(t) = \sum_{m \geq 0} m |\phi_m(t)|^2 = (O(t)|\hat{X}|O(t)), \quad \hat{X} = \sum_{m \geq 0} m |O_m\rangle\langle O_m|, \quad (5.4)$$

where $\hat{X}$ is the Krylov position (number) operator associated with the Krylov index. In this sense, $K(t) = \langle \hat{X} \rangle$ measures the mean distance, in Krylov space, of the evolving operator from the seed.

5.2.1 Continuous time: Lanczos

For Hamiltonian dynamics the Heisenberg evolution can be written as

$$|O(t)\rangle = e^{i\mathcal{L}t} |O_0\rangle, \quad \mathcal{L} = [\hat{H}, \cdot], \quad (5.5)$$

where $\mathcal{L}$ acts linearly on operator space. Starting from the normalised seed $|O_0\rangle$, define $|O_{-1}\rangle = 0$ and $b_0^{(\mathcal{L})} = 0$. The (operator) Lanczos recursion constructs an orthonormal Krylov basis $\{|O_m\rangle\}_{m \geq 0}$ together with recursion coefficients $\{a_m^{(\mathcal{L})}, b_{m \geq 1}^{(\mathcal{L})}\}$ via

$$\begin{aligned} a_m^{(\mathcal{L})} &= (O_m | \mathcal{L} | O_m), \\ |A_{m+1}\rangle &= \mathcal{L} | O_m \rangle - a_m^{(\mathcal{L})} | O_m \rangle - b_m^{(\mathcal{L})} | O_{m-1} \rangle, \\ b_{m+1}^{(\mathcal{L})} &= \| |A_{m+1}\rangle \|, \quad |O_{m+1}\rangle = |A_{m+1}\rangle / b_{m+1}^{(\mathcal{L})}. \end{aligned} \quad (5.6)$$

By construction, $|A_{m+1}\rangle$ is orthogonal to $|O_m\rangle$ and $|O_{m-1}\rangle$, and $b_{m+1}^{(\mathcal{L})} \geq 0$.

Equivalently, these coefficients are the matrix elements of the Liouvillian projected onto the Krylov subspace:

$$\tilde{\mathcal{L}}_{\ell m} = (O_\ell | \mathcal{L} | O_m) = a_m^{(\mathcal{L})} \delta_{\ell m} + b_{\ell+1}^{(\mathcal{L})} \delta_{\ell+1, m} + b_\ell^{(\mathcal{L})} \delta_{\ell-1, m}, \quad (5.7)$$

so that $\mathcal{L}$ is represented in the Krylov basis by the tridiagonal matrix

$$\tilde{\mathcal{L}}_{\ell m} = \begin{cases} a_m^{(\mathcal{L})}, & \ell = m, \\ b_m^{(\mathcal{L})}, & \ell = m \pm 1, \\ 0, & \ell \notin \{m-1, m, m+1\}, \end{cases} \quad (5.8)$$

If the initial seed is Hermitian the diagonal elements vanish,

$$a_m^{(\mathcal{L})} = (O_m | [\hat{H}, O_m]) = 0, \quad (5.9)$$

and the action reduces to the nearest-neighbour Krylov-chain form

$$\mathcal{L} | O_m \rangle = b_{m+1}^{(\mathcal{L})} | O_{m+1} \rangle + b_m^{(\mathcal{L})} | O_{m-1} \rangle, \quad b_0^{(\mathcal{L})} = 0. \quad (5.10)$$

Expanding the evolved operator as $|O(t)\rangle = \sum_{m \geq 0} \phi_m^{(\mathcal{L})}(t) |O_m\rangle$ defines amplitudes $\phi_m^{(\mathcal{L})}(t) = (O_m | O(t) \rangle)$ (with $\sum_m |\phi_m^{(\mathcal{L})}(t)|^2 = 1$). Projecting the evolution equation onto the Krylov basis yields a tight-binding dynamics on the half-chain,

$$i \partial_t \phi_m^{(\mathcal{L})}(t) = b_{m+1}^{(\mathcal{L})} \phi_{m+1}^{(\mathcal{L})}(t) + a_m^{(\mathcal{L})} \phi_m^{(\mathcal{L})}(t) + b_m^{(\mathcal{L})} \phi_{m-1}^{(\mathcal{L})}(t), \quad \phi_m^{(\mathcal{L})}(0) = \delta_{m0}, \quad \phi_{-1}^{(\mathcal{L})}(t) = 0, \quad (5.11)$$

(or the same equation without the on-site term when $a_m^{(\mathcal{L})} = 0$). The Krylov complexity $K(t)$ is given by the general definition above, here evaluated with the Lanczos amplitudes $\phi_m(t) = \phi_m^{(\mathcal{L})}(t)$.

Because $b_m^{(\mathcal{L})}$ are literally the off-diagonal entries of the tridiagonal $\tilde{\mathcal{L}}$ in Eq. Eq. (5.8), their large- m behaviour controls how fast the Krylov wavepacket moves away from the edge, and thus how rapidly complexity grows. $b_m^{(\mathcal{L})}$ appear as hopping matrix elements of the generator of motion on the Krylov chain, $\tilde{\mathcal{L}}$. This interpretation will not carry over verbatim to the discrete-time Arnoldi coefficients below.

5.2.2 Discrete time: Arnoldi

For unitary-circuit dynamics we instead evolve stroboscopically,

$$|O(t)\rangle = \mathcal{U}^t |O_0\rangle, \quad \mathcal{U}[\hat{O}] = \hat{U}^\dagger \hat{O} \hat{U}, \quad (5.12)$$

where $\mathcal{U}$ is unitary with respect to the inner product of Eq. (5.1). A direct Arnoldi iteration (Gram-Schmidt on $\{|O_0\rangle, \mathcal{U}|O_0\rangle, \mathcal{U}^2|O_0\rangle, \dots\}$) yields an orthonormal Krylov basis $\{|O_m\rangle\}_{m \geq 0}$ in which $\mathcal{U}$ is represented by an upper-Hessenberg matrix,

$$\tilde{\mathcal{U}}_{\ell m} = (O_\ell | \mathcal{U} | O_m), \quad \tilde{\mathcal{U}}_{\ell m} = 0 \quad \text{for } \ell > m + 1. \quad (5.13)$$

Equivalently, the Krylov amplitudes evolve as $\phi(t+1) = \tilde{\mathcal{U}} \phi(t)$, so $\tilde{\mathcal{U}}$ can be viewed as the one-step operator of a discrete-time quantum walk on the Krylov index. However, its matrix elements are overlap data: $\tilde{\mathcal{U}}_{\ell m}$ encodes the (gauge-dependent) projection of $\mathcal{U}|O_m\rangle$ onto $|O_\ell\rangle$, rather than a fixed “coupling” between Krylov sites m and ℓ (in contrast with the Lanczos chain for $\tilde{\mathcal{L}}$).

As for any orthonormal Krylov basis (including the Lanczos construction above), the vectors are only defined up to phases $|O_m\rangle \mapsto e^{i\theta_m} |O_m\rangle$, under which the matrix elements transform as

$$\tilde{\mathcal{U}}_{\ell m} \mapsto e^{-i\theta_\ell} e^{i\theta_m} \tilde{\mathcal{U}}_{\ell m}. \quad (5.14)$$

In the Hermitian Lanczos case this freedom is typically fixed implicitly by the recursion convention $b_{m+1}^{(\mathcal{L})} = \|A_{m+1}\| \geq 0$ (and, when desired, by choosing Hermitian Krylov operators). For unitary Arnoldi, by contrast, it is useful to keep track of this phase gauge: $\tilde{\mathcal{U}}$ is not tridiagonal, so there are many nonzero overlaps, and the choice of phases $\{\theta_m\}$ redistributes phases among these matrix elements (while leaving gauge-invariant quantities such as $|\phi_m(t)|^2$ unchanged). We first isolate the two pieces of Hessenberg data that have an invariant structural meaning, namely the diagonal and

subdiagonal entries,

$$a_m^{(\mathcal{U})} = \tilde{\mathcal{U}}_{mm} = (O_m | \mathcal{U} | O_m), \quad b_{m+1}^{(\mathcal{U})} = \tilde{\mathcal{U}}_{m+1,m} = (O_{m+1} | \mathcal{U} | O_m), \quad (m \geq 0). \quad (5.15)$$

The remaining phase freedom in Eq. (5.14) can be fixed recursively so that each subdiagonal entry is real and nonnegative,

$$b_m^{(\mathcal{U})} \in [0, 1], \quad m \geq 1, \quad (5.16)$$

which removes any ambiguity in the notion of “forward innovation” along the Krylov chain.¹

At this point the upper-triangular part of the Hessenberg matrix is still not fixed: it encodes how $\mathcal{U}|O_m\rangle$ projects back onto earlier Krylov vectors, and its detailed form depends on the chosen gauge. Rather than introducing additional coefficients directly from overlaps (which can obscure this gauge dependence), we define the remaining independent recursion data as the first-row entries of the gauge-fixed Hessenberg matrix,

$$c_m^{(\mathcal{U})} = \tilde{\mathcal{U}}_{0m}, \quad m \geq 1, \quad (5.17)$$

with $\tilde{\mathcal{U}}_{00} = a_0^{(\mathcal{U})}$ by definition.

Finally, exact unitarity of $\mathcal{U}$ imposes that all remaining matrix elements are fixed by $\{a_m^{(\mathcal{U})}, b_m^{(\mathcal{U})}, c_m^{(\mathcal{U})}\}$, yielding the canonical parameterisation

$$\tilde{\mathcal{U}}_{\ell m} = \begin{cases} 0, & \ell > m + 1, \\ b_{m+1}^{(\mathcal{U})}, & \ell = m + 1, \\ a_m^{(\mathcal{U})}, & \ell = m, \\ c_m^{(\mathcal{U})}, & \ell = 0, \quad m \geq 1, \\ \frac{a_\ell^{(\mathcal{U})}}{c_\ell^{(\mathcal{U})}} c_m^{(\mathcal{U})}, & 1 \leq \ell < m, \end{cases} \quad (5.18)$$

with the understanding that if some $c_\ell^{(\mathcal{U})} = 0$ the chain terminates (the Krylov subspace becomes finite-dimensional) and the above formula is applied only up to the termination index.

We stress that a finite $M \times M$ Hessenberg truncation obtained after M Arnoldi steps need not be unitary: unitarity holds for $\mathcal{U}$ on the full operator Hilbert space (and, equivalently, for the completed Hessenberg representation), whereas finite truncations

¹After enforcing $b_m^{(\mathcal{U})} \geq 0$ for all m , only an overall phase $\theta_m = \theta$ remains, which leaves $\tilde{\mathcal{U}}$ invariant and is therefore irrelevant.

are generically only contractive. The parameterisation Eq. (5.18) should therefore be interpreted as a statement about the exact (completed) unitary Hessenberg form associated with a cyclic seed and the gauge choice Eq. (5.16).

This makes two qualitative differences to continuous time explicit:

- The unitarity enforces $|\tilde{\mathcal{U}}_{\ell m}| \leq 1$ for all ℓ, m , hence $|a_m^{(\mathcal{U})}| \leq 1$ and $0 \leq b_m^{(\mathcal{U})} \leq 1$. The unbounded Lanczos mechanism $b_m^{(\mathcal{L})} \sim \alpha m$ responsible for sustained exponential $K_{\text{Lan}}(t)$ is therefore not available in genuinely unitary discrete-time dynamics.
- Compared to the Lanczos nearest-neighbour chain Eq. (5.6), unitary dynamics retains the forward “one-step” coupling $\tilde{\mathcal{U}}_{m+1,m} = b_{m+1}^{(\mathcal{U})}$ but generically also induces additional backward/long-range couplings through the upper-triangular structure Eq. (5.18), controlled by $a_m^{(\mathcal{U})}$ together with the (gauge-fixed) first-row data $c_m^{(\mathcal{U})}$ [307, 308, 310].

In discrete time the Krylov amplitudes have finite support at fixed stroboscopic time t . Since $\tilde{\mathcal{U}}$ is upper Hessenberg, one application of $\mathcal{U}$ can increase the Krylov index by at most one. Starting from the seed at $m = 0$, after t steps one therefore has $\phi_m^{(\mathcal{U})}(t) = 0$ for all $m > t$. The expansion thus reads

$$|O(t)\rangle = \sum_{m=0}^t \phi_m^{(\mathcal{U})}(t) |O_m\rangle. \quad (5.19)$$

As a consequence,

$$K_{\text{Arn}}(t) \leq t, \quad (5.20)$$

so linear growth is the largest possible asymptotic scaling in stroboscopic time.

It is conjectured [309] that in chaotic (and interacting integrable) circuits, the Krylov recursion data flows towards a maximally ergodic unitary (MEU) regime in which

$$a_m^{(\mathcal{U})} \rightarrow 0, \quad b_m^{(\mathcal{U})} \rightarrow 1 \quad (m \rightarrow \infty), \quad (5.21)$$

and the $a_m^{(\mathcal{U})}, c_m^{(\mathcal{U})}$ -controlled backward/long-range hopping terms vanish, leaving essentially unidirectional motion along the Krylov chain [309]. Equivalently, the Krylov basis operators become maximally ergodic: their autocorrelation functions vanish immediately and their spectral functions become flat on the Floquet Brillouin zone [309]. In this regime one finds asymptotically [309]

$$K_{\text{Arn}}(t) \sim t, \quad (5.22)$$

often with slope approaching the maximal value 1 (in units of ‘sites per Floquet step’).

For the remainder of [Part II](#), it is useful to treat the Lanczos and Floquet/Arnoldi constructions as the same conceptual object, a Krylov embedding of dynamics with complexity $K(t)$, while keeping in mind the key discrete-time constraint: unitarity bounds the recursion data and, in chaotic circuits, the flow is towards an MEU fixed point characterised by $b_m^{(\mathcal{U})} \rightarrow 1$ and $a_m^{(\mathcal{U})} \rightarrow 0$.

5.3 Connect Lanczos and Arnoldi

This section compares continuous-time Hamiltonian dynamics generated by a Liouvillian $\mathcal{L}$ and discrete-time Floquet dynamics generated by a unitary map $\mathcal{U}$ through spectral measures. Specifically, we consider the case when the discrete step is obtained by stroboscopic sampling (or ST) of an underlying $\mathcal{L}$,

$$\mathcal{U} = e^{-i\mathcal{L}\Delta t}. \quad (5.23)$$

Sampling at a finite step Δt induces an *aliasing* of frequencies: real-line spectral weight at ω and at $\omega + 2\pi k/\Delta t$ (for any $k \in \mathbb{Z}$) becomes indistinguishable in the stroboscopic data, because only the phase $\varphi = \omega\Delta t \pmod{2\pi}$ is retained.

For sufficiently small Δt one expects an intermediate-time window in which seed-resolved Krylov data and $K(t)$ resemble the Hamiltonian phenomenology, followed (at fixed $\Delta t > 0$) by a crossover to discrete-time behaviour at late stroboscopic times. In free or integrable circuit families this crossover can sharpen into a nonanalytic “Trotter transition” associated with structure near the Floquet π -mode².

The aim is twofold. First, we derive an exact “folding” relation between the real-line density $S_{\mathcal{L}}(\omega)$ and the unit-circle density $S_{\mathcal{U}}(\varphi)$, making precise how changing Δt reshapes the unit-circle measure and hence the unitary Hessenberg (Arnoldi) recursion of [Sec. 5.2](#). Second, we use two analytically controlled benchmarks to illustrate both the utility and the limitations of this spectral bridge.

In [Sec. 5.3.1](#) we start from a synthetic continuous-time Krylov model with constant Lanczos coefficients (a tridiagonal Toeplitz Liouvillian), compute $S_{\mathcal{L}}$ explicitly, and track how aliasing produces Δt -dependent unit-circle spectra $S_{\mathcal{U}}$ that interpolate between an arc-supported, semicircular-profile distribution (no aliasing) and an approximately flat distribution (strong aliasing).

In [Sec. 5.3.2](#) we address the inverse problem: starting from an MEU-type Hessenberg (flat $S_{\mathcal{U}}$), we show that $S_{\mathcal{U}}$ does *not* determine $S_{\mathcal{L}}$ uniquely because aliasing is many-to-one; even under the conservative principal-branch convention (single Nyquist interval),

²The Floquet π -mode is the zone-edge point $\varphi = \pi$ in the eigenphase spectrum of the one-step map $\mathcal{U}$. Spectral weight near $\varphi \simeq \pi$ yields an alternating component $s_m = \langle O_0 | \mathcal{U}^m | O_0 \rangle \sim (-1)^m$ and, via aliasing, collects contributions from real-line frequencies $\omega \simeq (\pi + 2\pi k)/\Delta t$.

the resulting continuous-time Lanczos coefficients are bounded and do not reproduce OGH-type growth.

This clarifies why there is no canonical one-to-one map between $\{b_m^{(\mathcal{L})}\}$ and $\{a_m^{(\mathcal{U})}, b_m^{(\mathcal{U})}, c_m^{(\mathcal{U})}\}$ without additional physical input (e.g. a choice of logarithm branch, bandwidth/locality constraints, or microscopic model information).

Define the seed autocorrelation function

$$C(t) = \text{Tr}[\hat{O}_0^\dagger \hat{O}_0(t)]/D = (O_0|e^{-i\mathcal{L}t}|O_0), \quad t \in \mathbb{R}, \quad (5.24)$$

and its spectral density on the real line by the Fourier pair³

$$S_{\mathcal{L}}(\omega) = \int_{-\infty}^{\infty} \frac{dt}{2\pi} e^{i\omega t} C(t), \quad C(t) = \int_{-\infty}^{\infty} d\omega e^{-i\omega t} S_{\mathcal{L}}(\omega). \quad (5.25)$$

Normalisation gives $\int d\omega S_{\mathcal{L}}(\omega) = C(0) = 1$.

Now fix $\Delta t > 0$ and form the stroboscopic unitary $\mathcal{U}$ in Eq. (5.23). Sampling $C(t)$ at integer multiples of Δt produces the discrete correlation sequence

$$s_m = (O_0|\mathcal{U}^m|O_0) = C(m\Delta t), \quad m \in \mathbb{Z}, \quad (5.26)$$

with associated unit-circle spectral density

$$S_{\mathcal{U}}(\varphi) = \frac{1}{2\pi} \sum_{m \in \mathbb{Z}} e^{i\varphi m} s_m, \quad s_m = \int_{-\pi}^{\pi} d\varphi e^{-i\varphi m} S_{\mathcal{U}}(\varphi), \quad \varphi \in [-\pi, \pi), \quad (5.27)$$

and $\int_{-\pi}^{\pi} d\varphi S_{\mathcal{U}}(\varphi) = s_0 = 1$.

Evaluating the inverse transform in Eq. (5.25) at $t = m\Delta t$ gives

$$s_m = \int_{-\infty}^{\infty} d\omega e^{-i\omega m\Delta t} S_{\mathcal{L}}(\omega). \quad (5.28)$$

Inserting Eq. (5.28) into Eq. (5.27) and using Poisson summation in the distributional sense [195],

$$\sum_{m \in \mathbb{Z}} e^{-imx} = 2\pi \sum_{k \in \mathbb{Z}} \delta(x - 2\pi k), \quad (5.29)$$

one obtains the exact “aliasing” identity

$$S_{\mathcal{U}}(\varphi) = \frac{1}{\Delta t} \sum_{k \in \mathbb{Z}} S_{\mathcal{L}}\left(\frac{\varphi + 2\pi k}{\Delta t}\right), \quad \varphi \in [-\pi, \pi). \quad (5.30)$$

³We use subscripts $\mathcal{L}$ and $\mathcal{U}$ to distinguish spectral densities associated with the Liouvillian and Floquet superoperator, respectively. This differs from the notation in Sec. 2.2, where subscripts denoted the operator. The seed (Hermitian) operator is held fixed throughout this Chapter.

Thus $S_{\mathcal{U}}$ is the Δt -dependent periodisation of $S_{\mathcal{L}}$ under the folding map $\varphi = \omega\Delta t \pmod{2\pi}$. In other words stroboscopic spectra forget where a frequency sits on the real line modulo $2\pi/\Delta t$.

The aliasing map [Eq. \(5.30\)](#) is the basic mechanism by which changing Δt can qualitatively change the discrete-time Hessenberg recursion discussed in [Sec. 5.2.2](#). Starting from a synthetic model already in a Krylov basis, we can build a direct pipeline to connect the Lanczos coefficients to the Arnoldi ones through [Eq. \(5.30\)](#)

$$\{b_m^{(\mathcal{L})}\} \rightarrow S_{\mathcal{L}} \xrightarrow[\text{sample at } \Delta t]{\text{folding}} S_{\mathcal{U}} \rightarrow \{a_m^{(\mathcal{U})}, b_m^{(\mathcal{U})}, c_m^{(\mathcal{U})}\}. \quad (5.31)$$

Three direct considerations are in order. First, one can take the two limits for Δt to 0 and ∞ .

(i) If the continuous spectral density is supported inside a single Nyquist interval,

$$S_{\mathcal{L}}(\omega) = 0 \quad \text{for } |\omega| \geq \pi/\Delta t, \quad (5.32)$$

then only the $k = 0$ term contributes in [Eq. \(5.30\)](#) and

$$S_{\mathcal{U}}(\varphi) = \frac{1}{\Delta t} S_{\mathcal{L}}\left(\frac{\varphi}{\Delta t}\right). \quad (5.33)$$

In this regime stroboscopic data contain (up to rescaling) the same seed-resolved spectral information as continuous time. This is the clean setting in which one expects an extended preasymptotic resemblance between continuous-time (Lanczos) and discrete-time (Arnoldi) Krylov dynamics. (ii) In the opposite case, without [Eq. \(5.32\)](#), different real-line spectra can induce the same unit-circle spectrum because frequencies separated by $2\pi/\Delta t$ are folded together. Moreover, if $S_{\mathcal{L}} \in L^1(\mathbb{R})$, then for fixed φ the right-hand side of [Eq. \(5.30\)](#) becomes a Riemann sum as $\Delta t \rightarrow \infty$, yielding a universal flattening:

$$S_{\mathcal{U}}(\varphi) \xrightarrow[\Delta t \rightarrow \infty]{} \frac{1}{2\pi}, \quad s_m \xrightarrow[\Delta t \rightarrow \infty]{} \delta_{m0}. \quad (5.34)$$

A flatter unit-circle spectral density corresponds to a more rapidly decorrelating stroboscopic sequence s_m and, through [Eq. \(5.31\)](#), to Hessenberg data closer to the late- m “maximally ergodic” fixed-point structure described in [Sec. 5.2.2](#) (characterised by $b_m^{(\mathcal{U})} \rightarrow 1$ and $a_m^{(\mathcal{U})} \rightarrow 0$). In this sense, aliasing provides a concrete spectral route by which increasing Δt (or, more generally, increasing the amount of folded high-frequency structure) pushes the discrete-time Krylov problem towards the regime in which $K_{\text{Arn}}(t)$ achieves its maximal linear growth.

(iii) Because [Eq. \(5.30\)](#) folds the real line onto $\varphi \in [-\pi, \pi)$, spectral structure near frequencies $\omega \simeq (\pi + 2\pi k)/\Delta t$ is precisely what controls $S_{\mathcal{U}}(\varphi)$ near $\varphi \simeq \pi$. This

is the natural spectral location where discrete-time effects that have no continuous-time counterpart can concentrate (e.g. suppression or gaps near $\varphi = \pi$ in integrable families), and it is the seed-resolved way to diagnose when the approach to the late- m MEU recursion pattern is obstructed or sharply modified. We will return to this viewpoint when interpreting how changing Δt can produce qualitatively different upper-Hessenberg structures even when the early-time dynamics remains close.

5.3.1 Constant Lanczos coefficients

A particularly transparent test case is the “free” Krylov chain in which the Lanczos recursion saturates immediately,

$$b_m^{(\mathcal{L})} = b > 0 \quad \text{for all } m \geq 1. \quad (5.35)$$

In the Krylov basis, the projected Liouvillian is then the semi-infinite tridiagonal Toeplitz (Jacobi) matrix

$$\tilde{\mathcal{L}} = b \sum_{m \geq 0} (|O_{m+1}\rangle\langle O_m| + |O_m\rangle\langle O_{m+1}|), \quad (5.36)$$

i.e. the tight-binding Hamiltonian on a half-line with constant hopping b and an open boundary at $m = 0$. This model can be treated analytically (both spectrally and in the time domain), providing a controlled setting in which to study the aliasing map [Eq. \(5.30\)](#) and its implications for the unitary Hessenberg data.

The eigenmodes of [Eq. \(5.36\)](#) are standing waves

$$\psi_q(m) = \sqrt{\frac{2}{\pi}} \sin[(m+1)q], \quad q \in [0, \pi], \quad (5.37)$$

with dispersion relation

$$\omega(q) = 2b \cos q, \quad \omega \in [-2b, 2b]. \quad (5.38)$$

For the edge seed $|O_0\rangle$, the spectral measure is weighted by $|\psi_q(0)|^2 = \frac{2}{\pi} \sin^2 q$. Pushing forward from q to ω therefore yields the seed-resolved spectral density on the real line,

$$S_{\mathcal{L}}(\omega) = \frac{1}{2\pi b^2} \sqrt{4b^2 - \omega^2} \mathbf{1}_{|\omega| \leq 2b}, \quad (5.39)$$

i.e. a normalised semicircle distribution supported on $[-2b, 2b]$.

Equivalently, the corresponding autocorrelation function is the Fourier transform of

Eq. (5.39),

$$C(t) = (O_0 | e^{-i\mathcal{L}t} | O_0) = \int_{-2b}^{2b} d\omega e^{-i\omega t} S_{\mathcal{L}}(\omega) = \frac{J_1(2bt)}{bt}, \quad (5.40)$$

with J_1 the Bessel function of the first kind. This provides a closed-form expression for the stroboscopic moments

$$s_m = (O_0 | \mathcal{U}^m | O_0) = C(m\Delta t) = \frac{J_1(2bm\Delta t)}{bm\Delta t}. \quad (5.41)$$

Applying the aliasing identity Eq. (5.30) to the compactly supported semicircle Eq. (5.39) gives an explicit unit-circle density as a finite sum,

$$S_{\mathcal{U}}(\varphi) = \frac{1}{2\pi b^2 \Delta t} \sum_{k \in \mathbb{Z}} \sqrt{4b^2 - \left(\frac{\varphi + 2\pi k}{\Delta t} \right)^2} \mathbf{1}_{|\varphi + 2\pi k| \leq 2b\Delta t}, \quad \varphi \in [-\pi, \pi). \quad (5.42)$$

Because $S_{\mathcal{L}}$ has support $|\omega| \leq 2b$, only those integers k for which $|\varphi + 2\pi k| \leq 2b\Delta t$ contribute, so the sum in Eq. (5.42) contains $O(b\Delta t/\pi)$ nonzero terms.

Two limiting regimes are particularly transparent.

- *No-aliasing (small step)*: if $2b \leq \pi/\Delta t$, i.e.

$$\Delta t \leq \frac{\pi}{2b}, \quad (5.43)$$

then the spectrum of $\mathcal{L}$ lies inside a single Nyquist interval and only $k = 0$ contributes in Eq. (5.42). One obtains

$$S_{\mathcal{U}}(\varphi) = \frac{1}{2\pi b^2 \Delta t} \sqrt{4b^2 - \left(\frac{\varphi}{\Delta t} \right)^2} \mathbf{1}_{|\varphi| \leq 2b\Delta t}. \quad (5.44)$$

Thus $S_{\mathcal{U}}$ is supported on a proper arc $|\varphi| \leq 2b\Delta t < \pi$, i.e. the Floquet spectrum exhibits a *gap* around the zone edge $\varphi = \pm\pi$.

- *Strong aliasing (large step)*: for $\Delta t > \pi/(2b)$, multiple translated semicircles overlap within $[-\pi, \pi)$ according to Eq. (5.42). As Δt increases, more copies contribute at each φ , and the distribution becomes progressively flatter. In the limit $\Delta t \rightarrow \infty$ one recovers the universal flattening Eq. (5.34).

The Nyquist threshold Eq. (5.43) marks a sharp change in *support*: for $\Delta t \leq \pi/(2b)$ the folded spectrum is arc-supported and gapped at the zone edge $\varphi = \pm\pi$, while for $\Delta t > \pi/(2b)$ overlap between translated copies fills the full Brillouin zone. Beyond this support effect, the subsequent approach of $S_{\mathcal{U}}$ towards the flat density $1/(2\pi)$ is a smooth aliasing-controlled crossover in this solvable model. If one nevertheless wants a scalar diagnostic of “how flat” the spectrum is (useful for tracking proximity to the

MEU limit), one can use the L^2 distance

$$\int_{-\pi}^{\pi} d\varphi \left| S_{\mathcal{U}}(\varphi) - \frac{1}{2\pi} \right|^2 = \frac{1}{2\pi} \sum_{m \neq 0} |s_m|^2 = \frac{1}{\pi} \sum_{m \geq 1} |s_m|^2, \quad s_m = C(m\Delta t), \quad (5.45)$$

where the second equality follows by Parseval's identity applied to [Eq. \(5.27\)](#).

5.3.2 MEU to Lanczos

This section addresses the “inverse” question to [Sec. 5.3](#): starting from a unitary Krylov problem in the MEU regime, what can one say about an underlying continuous-time generator $\mathcal{L}$ and its Lanczos coefficients $\{a_m^{(\mathcal{L})}, b_m^{(\mathcal{L})}\}$, and how do they depend on the stroboscopic step Δt ? The MEU fixed point corresponds to having constant $a_m^{(\mathcal{U})} = c_m^{(\mathcal{U})} = 0$ and $b_m^{(\mathcal{U})} = 1$, namely the Hessenberg is a unilateral shift, $\tilde{\mathcal{U}}_{\ell m} = \delta_{\ell, m+1}$. Equivalently, $\mathcal{U}|O_m\rangle = |O_{m+1}\rangle$ for all $m \geq 0$, so the corresponding moments are maximally decorrelated:

$$s_m = \langle O_0 | \mathcal{U}^m | O_0 \rangle = \delta_{m0}. \quad (5.46)$$

Using the Fourier pair [Eq. \(5.27\)](#), [Eq. \(5.46\)](#) leads to a flat unit-circle spectral density,

$$S_{\mathcal{U}}(\varphi) = \frac{1}{2\pi}, \quad \varphi \in [-\pi, \pi]. \quad (5.47)$$

Assume again $\mathcal{U} = e^{-i\mathcal{L}\Delta t}$. As discussed above, $S_{\mathcal{U}}$ does not uniquely determine $S_{\mathcal{L}}$. To remove aliasing ambiguity, the most conservative choice is the principal-branch assumption where all spectral weight lies inside a single Nyquist interval ($k = 0$), i.e. the bandlimit condition [Eq. \(5.32\)](#). In this case the map reduces to [Eq. \(5.33\)](#). Inserting the MEU density [Eq. \(5.47\)](#) into [Eq. \(5.33\)](#) gives

$$S_{\mathcal{L}}(\omega) = \frac{\Delta t}{2\pi} \mathbf{1}_{|\omega| \leq \pi/\Delta t}, \quad \int_{-\infty}^{\infty} d\omega S_{\mathcal{L}}(\omega) = 1. \quad (5.48)$$

Thus, on the principal branch, the inverse map sends a flat unit-circle spectrum to a *band-limited flat* real-line spectrum with bandwidth $\pi/\Delta t$.

The $b_m^{(\mathcal{L})}$ coefficients are those of rescaled Legendre polynomials (see e.g. [\[311\]](#) for the general orthogonal-polynomial Jacobi correspondence). Concretely, define $\frac{\pi}{\Delta t}$ to be the half-bandwidth and rescale $\omega = \pi/\Delta t x$ so that $x \in [-1, 1]$. With uniform weight on $[-1, 1]$, the orthonormal polynomials are proportional to Legendre polynomials, and their three-term recurrence yields

$$b_m^{(\mathcal{L})} = \frac{\pi}{\Delta t} \frac{m}{\sqrt{(2m-1)(2m+1)}}, \quad m \geq 1. \quad (5.49)$$

In particular,

$$b_1^{(\mathcal{L})} = \frac{\pi}{\sqrt{3} \Delta t}, \quad b_m^{(\mathcal{L})} \xrightarrow{m \rightarrow \infty} \frac{\pi}{2\Delta t}. \quad (5.50)$$

Under the principal-branch assumption, the generator reconstructed from an MEU step has *bounded* Lanczos coefficients: the associated tridiagonal Jacobi matrix is asymptotically Toeplitz (constant off-diagonal $\pi/(2\Delta t)$). This within the OGH corresponds to the growth regime associated with non interacting models. At the contrary, $b_m^{(\mathcal{L})} \sim \alpha m$, and it implies that any sustained exponential growth of continuous-time Krylov complexity is absent in this reconstructed model.

Aliasing implies that one can engineer infinitely many distinct densities $S_{\mathcal{L}}$ whose folding under $\varphi = \omega \Delta t \pmod{2\pi}$ yields the same flat $S_{\mathcal{U}}(\varphi) = 1/(2\pi)$. Such choices generically modify the Lanczos coefficients and can produce qualitatively different $\{b_m^{(\mathcal{L})}\}$, including sequences that grow over some range.

In summary, an MEU Hessenberg determines, under principal-branch convention, a band-limited flat density Eq. (5.48) and the corresponding bounded Lanczos coefficients Eq. (5.49), scaling as $1/\Delta t$ and saturating to $\pi/(2\Delta t)$ at large Krylov depth.

Proof: *Lanczos coefficients for a flat band-limited density.*

In the principal-branch reconstruction one takes

$$S_{\mathcal{L}}(\omega) = \frac{\Delta t}{2\pi} \mathbf{1}_{|\omega| \leq \pi/\Delta t} = \frac{1}{2A} \mathbf{1}_{|\omega| \leq A}, \quad A = \frac{\pi}{\Delta t}, \quad (5.51)$$

and defines the probability measure $d\mu(\omega) = S_{\mathcal{L}}(\omega) d\omega$. Let $\{p_m(\omega)\}$ be orthonormal polynomials for $d\mu$. Multiplication by ω has the Jacobi form

$$\omega p_m(\omega) = b_{m+1}^{(\mathcal{L})} p_{m+1}(\omega) + a_m^{(\mathcal{L})} p_m(\omega) + b_m^{(\mathcal{L})} p_{m-1}(\omega), \quad (5.52)$$

with $p_{-1} = 0$. Since $d\mu$ is even, $a_m^{(\mathcal{L})} = \int \omega p_m(\omega)^2 d\mu(\omega) = 0$.

Rescale to $x = \omega/A \in [-1, 1]$. Then

$$d\mu(\omega) = \frac{1}{2A} d\omega = \frac{1}{2} dx = d\nu(x), \quad (5.53)$$

so the orthonormal polynomials for $d\nu$ are proportional to Legendre polynomials.

Using $\int_{-1}^1 P_m(x) P_n(x) dx = \frac{2}{2m+1} \delta_{mn}$, an orthonormal choice for $d\nu = \frac{1}{2} dx$ is

$$q_m(x) = \sqrt{2m+1} P_m(x). \quad (5.54)$$

The standard Legendre recurrence

$$x P_m(x) = \frac{m+1}{2m+1} P_{m+1}(x) + \frac{m}{2m+1} P_{m-1}(x) \quad (5.55)$$

then implies, in the orthonormal basis $\{q_m\}$,

$$x q_m(x) = \frac{m+1}{\sqrt{(2m+1)(2m+3)}} q_{m+1}(x) + \frac{m}{\sqrt{(2m-1)(2m+1)}} q_{m-1}(x). \quad (5.56)$$

Hence the off-diagonal Jacobi coefficients for multiplication by x are

$$b_m^{(x)} = \frac{m}{\sqrt{(2m-1)(2m+1)}}, \quad m \geq 1. \quad (5.57)$$

Finally, since $\omega = Ax$, multiplication by ω scales the Jacobi coefficients by A , i.e.

$$b_m^{(\mathcal{L})} = A b_m^{(x)} = \frac{\pi}{\Delta t} \frac{m}{\sqrt{(2m-1)(2m+1)}}, \quad m \geq 1, \quad (5.58)$$

which is [Eq. \(5.49\)](#).

The aliasing identity [Eq. \(5.30\)](#) gives an exact map from the real-line seed-resolved spectral density $S_{\mathcal{L}}$ to the unit-circle density $S_{\mathcal{U}}$ obtained by stroboscopic sampling at step Δt , and hence (via the moment problem) to the corresponding unitary Hessenberg recursion. However, the folding it describes, $\omega \mapsto \varphi = \omega \Delta t \pmod{2\pi}$, is many-to-one, so relating $\tilde{\mathcal{L}}$ (and its Lanczos coefficients) to $\tilde{\mathcal{U}}$ (and its Arnoldi/Hessenberg data) is not well-posed without additional input such as a branch choice for the logarithm and/or constraints on the admissible real-line spectrum. The examples seen in [Secs. 5.3.1](#) and [5.3.2](#), make the limitation of the spectral bridge concrete: without extra structure beyond the folding map, the Arnoldi coefficients seem not to be able to be connected to specific OGH-type models.

This leads to the notorious problem of finding the correct Floquet Hamiltonian, which has been studied in various contexts [[23](#), [24](#), [99](#), [312](#)]. Our final goal will be to connect Floquet systems with properties unique to discrete-time dynamics, such as DU circuits, to continuous-time Hamiltonian systems in an adiabatic way. Fundamental in this direction has been the work of Schindler and Bukov, who constructed a family of Floquet Hamiltonians that interpolate between discrete-time and continuous-time dynamics [[313](#)]. In the next sections we introduce the Floquet theorem, which will be the main tool to construct such an interpolation.

5.4 Floquet Hamiltonians

Discrete-time dynamics generated by periodic driving can often be described in terms of an effective time-independent Hamiltonian, the *Floquet Hamiltonian*. This section introduces the Floquet Hamiltonian and the Floquet theorem, which provides a structural decomposition of the time-evolution operator for periodically driven systems.

We then survey common constructions of $\hat{H}_F$, emphasising the intrinsic gauge/branch ambiguities. Finally, we use the driven XX chain as a controlled benchmark for these issues and motivate the MEO principle, which will be introduced in the next section.

Let $\hat{H}(t)$ be τ -periodic, $\hat{H}(t + \tau) = \hat{H}(t)$, and define the propagator

$$\hat{U}(t, 0) = \mathcal{T} \exp\left(-i \int_0^t dt' \hat{H}(t')\right), \quad \hat{U} = \hat{U}(\tau, 0). \quad (5.59)$$

It is common to introduce an effective time-independent Floquet Hamiltonian $\hat{H}_F$ by

$$\hat{U} = e^{-i\tau \hat{H}_F}. \quad (5.60)$$

Eq. (5.60) is a *choice of logarithm* of $\hat{U}$; hence $\hat{H}_F$ is not unique. At the spectral level, quasienergies are defined modulo $\omega = 2\pi/\tau$, and different branch choices (or, equivalently, different Floquet gauge conventions) can yield distinct $\hat{H}_F$ that generate the same stroboscopic unitary $\hat{U}$. A more precise structural statement is the Floquet theorem [23, 24, 99, 314–316]: there exists a *periodic* unitary micromotion operator $\hat{P}(t)$, with $\hat{P}(t + \tau) = \hat{P}(t)$, and a *time-independent* Hermitian operator $\hat{H}_F$ such that

$$\hat{U}(t, 0) = \hat{P}(t) e^{-it \hat{H}_F} \hat{P}^\dagger(0). \quad (5.61)$$

In particular, at stroboscopic times $t = k\tau$ one has $\hat{P}(k\tau) = \hat{P}(0)$ and therefore

$$\hat{U}(k\tau, 0) = \hat{P}(0) e^{-ik\tau \hat{H}_F} \hat{P}^\dagger(0), \quad \hat{U} = \hat{P}(0) e^{-i\tau \hat{H}_F} \hat{P}^\dagger(0). \quad (5.62)$$

A frequently used Floquet gauge is $\hat{P}(0) = \hat{I}$, in which case Eq. (5.62) reduces to $\hat{U} = e^{-i\tau \hat{H}_F}$, consistent with Eq. (5.60).

It is often useful (and fully equivalent) to rewrite Floquet's theorem in the *rotating-frame* form used throughout the modern Floquet literature. Define the co-moving propagator

$$\hat{U}_{\text{rot}}(t) = \hat{P}^\dagger(t) \hat{U}(t, 0) \hat{P}(0). \quad (5.63)$$

Differentiating and using $i\partial_t \hat{U}(t, 0) = \hat{H}(t) \hat{U}(t, 0)$ gives

$$i\partial_t \hat{U}_{\text{rot}}(t) = \left[\hat{P}^\dagger(t) \hat{H}(t) \hat{P}(t) - i\hat{P}^\dagger(t) \partial_t \hat{P}(t) \right] \hat{U}_{\text{rot}}(t). \quad (5.64)$$

Floquet's theorem is precisely the statement that one can choose $\hat{P}(t)$ such that the bracket in Eq. (5.64) is *time independent*; identifying this constant with $\hat{H}_F$ yields the rotating-frame identity

$$\hat{H}_F = \hat{P}^\dagger(t) \hat{H}(t) \hat{P}(t) - i\hat{P}^\dagger(t) \partial_t \hat{P}(t), \quad (5.65)$$

and then $\hat{U}_{\text{rot}}(t) = e^{-it\hat{H}_F}$ implies Eq. (5.61). Conversely, inserting Eq. (5.61) into the Schrödinger equation reproduces Eq. (5.65). Thus Eq. (5.61) and Eq. (5.65) are the same theorem expressed in two equivalent languages (micromotion decomposition versus static generator in a periodic rotating frame).

The pair $(\hat{P}(t), \hat{H}_F)$ is not unique. Beyond the logarithm/branch ambiguity already present in Eq. (5.60), one may perform time-periodic gauge transformations that reshuffle phase between micromotion and the static generator while leaving $\hat{U}$ invariant. In practice, different constructions of $\hat{H}_F$ correspond to different (often approximate) realisations of this Floquet gauge freedom: they agree on $\hat{U}$ but can differ substantially at the level of the chosen generator and its real-space structure.

Methods to construct $\hat{H}_F$ can be grouped as follows [23, 24, 99]: (i) *exact logarithm* $\hat{H}_F = (i/\tau) \log \hat{U}$ with a specified branch; (ii) *high-frequency* ($1/\omega$) expansions that yield $\hat{H}_F$ and $\hat{P}(t)$ order by order, including Floquet-Magnus [317, 318], van Vleck (block-diagonalisation in Sambe space) [319, 320], and Brillouin-Wigner (projective perturbation theory in Sambe space) [321, 322]; (iii) *product-formula* settings (Trotter layers) where $\hat{H}_F$ is expanded via BCH series for $\log \prod_\ell e^{-i\Delta t_\ell \hat{H}_\ell}$ [323]; (iv) *resonant* regimes treated by (extended) degenerate perturbation theory in Floquet-Hilbert (Sambe) space to obtain effective generators on enlarged near-degenerate subspaces [324]; and (v) *geometric and unwinding* approaches that select a representative by additional structure, such as phase unwinding in Sambe/Fourier space [325] or geometric Floquet theory separating one-cycle evolution into holonomy and a dynamical generator (average-energy operator) [313].

In the next sections we introduce the MEO method, which provides a physically motivated principle to select among the many possible Floquet Hamiltonians by seeking a continuous unfolding of eigenphases and eigenvectors that remains approximately local in real space along an adiabatic path from continuous-time to discrete-time dynamics in Sec. 5.5, then apply it to the driven XX chain in Sec. 5.6 which offers an exact benchmark. Moreover, we build an adiabatic path whose endpoint is a circuit composed of DU gates. According to the transfer-matrix characterisation of Ref. [41], dual-unitary circuits can realise qualitatively distinct dynamical behaviours, ranging from non-interacting and non-ergodic (integrable) cases to ergodic but non-mixing, and fully ergodic and mixing regimes. They are the natural “ultimate” discrete-time models, as they realise extremal light-cone propagation and entangling dynamics.

5.5 Maximum Eigenprojector Overlap

The preceding discussion highlighted that different Floquet gauge choices yield different generators $\hat{H}_F$ for the same stroboscopic map $\hat{U}$. Among the infinitely many possibilities, we now seek a *continuous* unfolding that remains approximately local in real space along a chosen *path in unitary space*.

Throughout, the word “adiabatic” refers *only* to a smooth continuation in a control parameter: we consider a curve $\tau \mapsto \hat{U}(\tau)$ in the unitary group, and we use it as a mathematical device to connect and label eigenprojectors between nearby points on the curve. Importantly, $\hat{U}(\tau)$ is *not* the propagator of some intermediate-time evolution within a single experiment; rather, for each fixed τ it denotes a distinct *one-cycle* (stroboscopic) unitary of a different drive. The continuation problem is therefore a *branch-selection* problem for the matrix logarithm along the curve.

5.5.1 Floquet path and tracking

We consider a two-step brickwork protocol with fixed Hamiltonians $\hat{H}_1$ and $\hat{H}_2$, and use the half-step strength $\tau \in [0, 1]$ as the path parameter. For each τ , the *one-cycle* unitary is

$$\hat{U}(\tau) = e^{-i\tau \hat{H}_2} e^{-i\tau \hat{H}_1}, \quad \tau \in [0, 1]. \quad (5.66)$$

This is the Suzuki-Trotter (ST) path: $\hat{H}_1$ and $\hat{H}_2$ remain fixed while the gate strength τ is increased. More general paths where $\hat{H}_{1,2}$ vary with τ can be treated analogously; we focus on ST for concreteness.

For numerical tracking we discretise the path,

$$\tau_j = \tau_{\min} + j \Delta\tau, \quad j = 0, 1, \dots, M_\tau, \quad \Delta\tau = \frac{1 - \tau_{\min}}{M_\tau}, \quad (5.67)$$

so the algorithm acts on the sequence $\{\hat{U}(\tau_j)\}_{j=0}^{M_\tau}$ by matching spectral data between successive points along the path.

We define the Floquet generator via a branch choice of the logarithm,

$$\hat{U}(\tau) = e^{-i\tau \hat{H}_F(\tau)}, \quad \hat{H}_F(\tau) = \frac{i}{\tau} \log \hat{U}(\tau). \quad (5.68)$$

In the small- τ limit, the Baker-Campbell-Hausdorff expansion gives

$$\hat{U}(\tau) \approx e^{-i\tau(\hat{H}_1 + \hat{H}_2) + \mathcal{O}(\tau^2)}.$$

Thus $\hat{H}_F(\tau \rightarrow 0) \rightarrow \hat{H}_1 + \hat{H}_2$.

At fixed τ , the eigenvalues of $\hat{U}(\tau)$ satisfy $\mu_i(\tau) = e^{-i\varphi_i(\tau)}$ where $\varphi_i(\tau) \in (-\pi, \pi]$ are the *principal eigenphases*.

From these, we define several related quantities:

- The *(folded) quasienergies* (principal branch): $\epsilon_i(\tau) = \frac{\varphi_i(\tau)}{\tau}$. These take values in $(-\pi/\tau, \pi/\tau]$.
- The *unfolded eigenphases*: $\tilde{\varphi}_i(\tau) \in \mathbb{R}$, obtained by enforcing continuity along the τ -path, $\tilde{\varphi}_i(\tau) = \varphi_i(\tau) + 2\pi k_i(\tau)$ with integers $k_i(\tau)$ chosen to minimise discontinuities.
- The *unfolded quasienergies*: $\tilde{\epsilon}_i(\tau) = \tilde{\varphi}_i(\tau)/\tau \in \mathbb{R}$, which the Floquet generator $\hat{H}_F(\tau)$ uses as its eigenvalues.
- The *continuous-time limit energies*: $E_i = \lim_{\tau \rightarrow 0} \tilde{\epsilon}_i(\tau)$, which correspond to the eigenvalues of the continuous Hamiltonian $\hat{H}_1 + \hat{H}_2$.

These unfolded quantities are defined by enforcing continuity along the discretised path, up to unavoidable ambiguities at degeneracies and crossings addressed below.

5.5.2 Eigenprojectors near degeneracies

Near degeneracies, individual eigenvectors can rotate within the degenerate subspace and are therefore numerically unstable. By contrast, the projector onto the subspace is unique and can be tracked reliably. We compute a unitary Schur decomposition

$$\hat{U}(\tau) = \hat{V}(\tau) \hat{T}(\tau) \hat{V}^\dagger(\tau), \quad (5.69)$$

where $\hat{V}$ is unitary and $\hat{T}$ is upper-triangular with diagonal $\mu_i(\tau)$. We use the Schur form because it preserves unitarity (well-conditioned columns), exposes eigenvalues on the diagonal without a fragile diagonalisation, and groups nearly degenerate modes into a controllable triangular block so that projectors remain stable even when individual Schur vectors ($|v_i(\tau)\rangle$, columns of $\hat{V}(\tau)$) rotate. Using a tolerance δ_μ , we cluster nearby eigenvalues on the unit circle to form blocks $\mathcal{B}_a(\tau)$ with projectors

$$\hat{P}_a(\tau) = \sum_{i \in \mathcal{B}_a(\tau)} |v_i(\tau)\rangle \langle v_i(\tau)|, \quad r_a(\tau) = |\mathcal{B}_a(\tau)|. \quad (5.70)$$

As τ varies, blocks may split or merge depending on the local spectral geometry. Allowing this behaviour is essential for robust tracking through degeneracies.

5.5.3 Assignment via projector overlap

Given a set of tracked projectors $\{\hat{P}_a(\tau_j)\}$ and new projectors $\{\hat{P}'_b(\tau_{j+1})\}$, we quantify their match by

$$s_{ab} = \frac{1}{r_a} \text{Tr}[\hat{P}_a(\tau_j) \hat{P}'_b(\tau_{j+1})] = \frac{1}{r_a} \|\hat{B}_a^\dagger \hat{B}'_b\|_F^2, \quad (5.71)$$

where $\hat{B}_a$ and $\hat{B}'_b$ are orthonormal bases for the two subspaces. This definition remains meaningful even when ranks differ, which is crucial when blocks split or merge. To obtain a permutation of individual Schur vectors, we lift the block overlaps to a $D \times D$ score by replicating ranks:

$$\Sigma_{ij} = s_{a(i), b(j)}. \quad (5.72)$$

We then determine the permutation π maximising $\sum_i \Sigma_{i, \pi(i)}$ via the Hungarian method (maximum-weight bipartite matching) [326, 327], which yields a robust coarse match through degeneracies and block reshuffling.

5.5.4 Quasienergy unfolding

After tracking the evolution of the eigenprojectors, we assign logarithm branches so that the unfolded phases $\tilde{\varphi}_i(\tau_{j+1})$ vary continuously along τ . With π fixed, we unfold the principal phases against a predictor $\tilde{\varphi}_i(\tau_j)$:

$$\begin{aligned} \tilde{\epsilon}_i(\tau_{j+1}) &= \frac{\tilde{\varphi}_i(\tau_{j+1})}{\tau_{j+1}} = \frac{\varphi_i(\tau_{j+1}) + 2\pi k_i(\tau_{j+1})}{\tau_{j+1}}, \\ k_i(\tau_{j+1}) &= \text{round}\left(\frac{\tilde{\varphi}_i(\tau_j) - \varphi_i(\tau_{j+1})}{2\pi}\right). \end{aligned} \quad (5.73)$$

The predictor is obtained by short-step extrapolation of the previously accepted unfolded phases, or from the continuous-time baseline at the earliest steps. This enforces a continuous real-valued branch for $\tilde{\varphi}_i(\tau_{j+1})$ (and hence for $\tilde{\epsilon}_i(\tau_{j+1})$) along the discretised path.

Given the tracked Schur basis $\{|v_i(\tau_{j+1})\rangle\}$ and the unfolded phases $\{\tilde{\varphi}_i(\tau_{j+1})\}$, we construct the unfolded generator via the spectral decomposition

$$\hat{H}_F(\tau_{j+1}) = \sum_i \tilde{\epsilon}_i(\tau_{j+1}) |v_i(\tau_{j+1})\rangle \langle v_i(\tau_{j+1})|, \quad (5.74)$$

and verify numerically that $e^{-i\tau_{j+1}\hat{H}_F(\tau_{j+1})} \approx \hat{U}(\tau_{j+1})$ to the chosen tolerance.

5.5.5 Locality tie-breaker

When the overlap scores for candidate permutations are nearly tied, overlap alone cannot decide between crossing and repulsion over a single step. In this case, we choose the labelling that minimises the real-space nonlocality of $\hat{H}_F$. Concretely, we compute a locality functional and use the normalised average Pauli-string weight

$$\text{Loc}(\hat{H}) = \frac{\sum_P w(P) |c_P|^2}{\sum_P |c_P|^2}, \quad w(P) = N - \#\{\text{I in } P\}, \quad (5.75)$$

where $\hat{H} = \sum_P c_P P$ is the Pauli expansion on N sites. We then set

$$\pi_\star = \arg \min_{\pi \in \mathcal{C}} \text{Loc}(\hat{H}_F[\pi]), \quad (5.76)$$

where $\mathcal{C}$ contains local swaps among nearby modes (close in phase and nearly tied in overlap), which keeps the search small. Mislabeling a narrow avoided crossing typically induces discontinuities in the unfolded phases and long-range terms in $\hat{H}_F$; by contrast, the correct continuation tends to preserve locality. Genuine symmetry-decoupled crossings are naturally allowed by the criterion.

This locality preference is consistent with quasi-adiabatic continuation and with adiabatic gauge potentials, both of which favour locality-preserving gauges for generators [328, 329].

5.6 XX spin chain under brickwork drive

Having introduced the MEO method, we benchmark it on a free-fermion setting where the Floquet problem admits an exact single-particle reduction: the brickwork-driven XX chain along a period-scaling path $\tau \in [0, 1]$ connecting the continuous-time limit $\tau \rightarrow 0$ to the DU endpoint $\tau = 1$. After Jordan-Wigner, the dynamics is quadratic and translation invariant with a two-site unit cell, so the Floquet unitary factorises into independent 2×2 reduced-momentum blocks. This yields closed-form eigenphases and, importantly, a fully analytic unfolding prescription: from the trace one obtains the principal (folded) phase magnitude $\varphi(k; \tau) \in [0, \pi]$, and the missing sign is fixed *without any path tracking* by matching to the $\tau \rightarrow 0$ continuous-time dispersion, i.e. by setting $\tilde{\vartheta}(k; \tau) = \text{sgn}(\cos(k/2)) \varphi(k; \tau)$. This provides an unambiguous reference for what is meant by *folded* versus *unfolded* quasienergies, against which MEO can be tested [312, 313].

In what follows, we first define the model and its single-particle reduction, and record the exact folded phases. We then construct unfolded single-particle quasienergies via

the analytic sign convention above, and build the many-body spectrum by filling modes (without any mod 2π reduction after the sum). We compare the resulting analytical Floquet generators with those reconstructed by MEO, and finally discuss the Krylov complexity at the DU endpoint.

We consider an even number N of spins with periodic boundary conditions and the homogeneous XX Hamiltonian

$$\hat{H}_{\text{XX}} = \frac{\pi}{4} \sum_{n=1}^N \left(\hat{\sigma}_n^x \hat{\sigma}_{n+1}^x + \hat{\sigma}_n^y \hat{\sigma}_{n+1}^y \right), \quad (n + N = n). \quad (5.77)$$

Split it into commuting bond layers on odd and even links,

$$\hat{H}_{\text{odd}} = \frac{\pi}{4} \sum_{n=1}^{N/2} \left(\hat{\sigma}_{2n-1}^x \hat{\sigma}_{2n}^x + \hat{\sigma}_{2n-1}^y \hat{\sigma}_{2n}^y \right), \quad \hat{H}_{\text{even}} = \frac{\pi}{4} \sum_{n=1}^{N/2} \left(\hat{\sigma}_{2n}^x \hat{\sigma}_{2n+1}^x + \hat{\sigma}_{2n}^y \hat{\sigma}_{2n+1}^y \right), \quad (5.78)$$

so that $\hat{H}_{\text{XX}} = \hat{H}_{\text{odd}} + \hat{H}_{\text{even}}$. We define a two-step brickwork cycle of duration τ by the piecewise-constant drive

$$\hat{H}(t) = \begin{cases} 2\hat{H}_{\text{even}}, & 0 \leq t < \tau/2, \\ 2\hat{H}_{\text{odd}}, & \tau/2 \leq t < \tau, \end{cases} \quad \implies \quad \hat{U}(\tau) = e^{-i\tau\hat{H}_{\text{odd}}} e^{-i\tau\hat{H}_{\text{even}}}. \quad (5.79)$$

The Floquet generator is defined by

$$\hat{U}(\tau) = e^{-i\tau\hat{H}_F(\tau)}, \quad \hat{H}_F(\tau) = \frac{i}{\tau} \log \hat{U}(\tau), \quad (5.80)$$

where log denotes a branch choice. By a BCH expansion,

$$\hat{U}(\tau) = \exp \left[-i\tau(\hat{H}_{\text{odd}} + \hat{H}_{\text{even}}) + \mathcal{O}(\tau^2) \right], \quad \implies \quad \hat{H}_F(\tau \rightarrow 0) = \hat{H}_{\text{XX}}. \quad (5.81)$$

At $\tau = 1$ each half-cycle applies the two-site gate $e^{-i\frac{\pi}{4}(\hat{\sigma}^x \otimes \hat{\sigma}^x + \hat{\sigma}^y \otimes \hat{\sigma}^y)}$ on alternating bonds; this gate is DU, hence $\tau = 1$ is the DU endpoint.

Under Jordan-Wigner the drive is number conserving and quadratic,

$$\hat{H}_{\text{layer}} = \sum_{n,m=1}^N h_{nm}^{(\text{layer})} \hat{c}_n^\dagger \hat{c}_m, \quad \text{layer} \in \{\text{odd}, \text{even}\}, \quad (5.82)$$

so $\hat{U}(\tau)$ is Gaussian and acts linearly on fermions:

$$\hat{\mathbf{c}} \mapsto \hat{\mathbf{u}}(\tau) \hat{\mathbf{c}}, \quad \hat{\mathbf{u}}(\tau) = e^{-i\tau\hat{h}^{(\text{odd})}} e^{-i\tau\hat{h}^{(\text{even})}}.$$

Using the odd/even partition as a two-site unit cell, write $N = 2L$ with sublattices

$A = 2n - 1$ and $B = 2n$ and define reduced Fourier modes

$$\hat{\psi}_k = \begin{pmatrix} \hat{c}_{k,A} \\ \hat{c}_{k,B} \end{pmatrix}, \quad \hat{c}_{k,A} = \frac{1}{\sqrt{L}} \sum_{n=1}^L e^{-ikn} \hat{c}_{2n-1}, \quad \hat{c}_{k,B} = \frac{1}{\sqrt{L}} \sum_{n=1}^L e^{-ikn} \hat{c}_{2n}. \quad (5.83)$$

In a fixed fermion-parity sector $\hat{P} = \pm 1$ (spin PBC implies parity-dependent fermionic boundary conditions) the allowed reduced momenta are

$$K_+ = \left\{ \frac{(2m+1)\pi}{L} : m = 0, \dots, L-1 \right\}, \quad K_- = \left\{ \frac{2m\pi}{L} : m = 0, \dots, L-1 \right\}. \quad (5.84)$$

Accordingly $\hat{u}(\tau) = \bigoplus_{k \in K_{\hat{P}}} \hat{u}_k(\tau)$, and in the basis [Eq. \(5.83\)](#) the layer generators take the Pauli form

$$\hat{h}_k^{(\text{odd})} = \frac{\pi}{2} \hat{\sigma}^x, \quad \hat{h}_k^{(\text{even})} = \frac{\pi}{2} (\cos k \hat{\sigma}^x + \sin k \hat{\sigma}^y), \quad (5.85)$$

so that

$$\hat{u}_k(\tau) = e^{-i\tau \hat{h}_k^{(\text{odd})}} e^{-i\tau \hat{h}_k^{(\text{even})}} = e^{-i\frac{\pi}{2}\tau \hat{\sigma}^x} e^{-i\frac{\pi}{2}\tau (\cos k \hat{\sigma}^x + \sin k \hat{\sigma}^y)}. \quad (5.86)$$

5.6.1 Analytic phases and unfolding

For each reduced momentum k , the 2×2 block $\hat{u}_k(\tau) \in SU(2)$ has eigenvalues $e^{-i\vartheta_{k,+}(\tau)}$ and $e^{-i\vartheta_{k,-}(\tau)}$, where each eigenphase $\vartheta_{k,a}$ is only defined modulo 2π . A standard *folded* representative is obtained by fixing the *magnitude* $\varphi(k; \tau) \in [0, \pi]$ through the trace:

$$\cos \varphi(k; \tau) = \frac{1}{2} \text{Tr}[\hat{u}_k(\tau)] = \cos^2\left(\frac{\pi}{2}\tau\right) - \sin^2\left(\frac{\pi}{2}\tau\right) \cos k. \quad (5.87)$$

Because the right-hand side of [Eq. \(5.87\)](#) lies in $[-1, 1]$, the principal arccos is always well-defined. However, φ is *folded by construction*: it discards the sign information contained in the two eigenphases $\pm\vartheta \pmod{2\pi}$, retaining only $|\vartheta|$ in $[0, \pi]$.

To recover an *unfolded* (signed) eigenphase without any phase-tracking along the τ -path, we fix the sign by matching to the continuous-time limit $\tau \rightarrow 0$, where $\hat{H}_F(\tau) \rightarrow \hat{H}_{XX}$ and the single-particle dispersion in the reduced Brillouin zone is proportional to $\cos(k/2)$. Expanding [Eq. \(5.87\)](#) at small τ gives

$$\varphi(k; \tau) = \pi\tau |\cos(k/2)| + \mathcal{O}(\tau^3), \quad (5.88)$$

i.e. the principal convention naturally produces the *absolute value* of the desired

linear-in- τ phase. We therefore define the *unfolded signed phase* by the analytic rule

$$\tilde{\vartheta}(k; \tau) = s(k) \varphi(k; \tau) \in (-\pi, \pi], \quad s(k) = \text{sgn}(\cos(k/2)), \quad (5.89)$$

(with the harmless convention $s(k) = +1$ when $\cos(k/2) = 0$). Then the two unfolded eigenphases are $\tilde{\vartheta}_{k,\pm}(\tau) = \pm \tilde{\vartheta}(k; \tau)$ (now viewed as real numbers, not reduced modulo 2π), and the corresponding unfolded quasienergies (per cycle of duration τ) are

$$\tilde{\epsilon}_{k,\pm}(\tau) = \frac{\tilde{\vartheta}_{k,\pm}(\tau)}{\tau} = \pm \frac{\tilde{\vartheta}(k; \tau)}{\tau} \in \mathbb{R}. \quad (5.90)$$

By construction, [Eq. \(5.89\)](#) ensures the correct continuous-time limit

$$\tilde{\epsilon}_{k,\pm}(\tau \rightarrow 0) = \pm \pi \cos(k/2), \quad (5.91)$$

which matches the single-particle spectrum of $\hat{H}_{\text{XX}}$ in the reduced-zone parametrisation.

Since the drive is number conserving and Gaussian, many-body Floquet eigenstates are obtained by occupying single-particle Floquet modes. In a fixed fermion-parity sector $\hat{P} = \pm 1$ (hence $k \in K_{\hat{P}}$), the unfolded many-body quasienergy is the additive subset sum

$$\tilde{\epsilon}_{\{n\}}(\tau) = \sum_{k \in K_{\hat{P}}} \sum_{a=\pm} n_{k,a} \tilde{\epsilon}_{k,a}(\tau), \quad n_{k,a} \in \{0, 1\}, \quad (5.92)$$

and, crucially, the result of [Eq. \(5.92\)](#) is *not* reduced modulo 2π . Here $\{n\}$ denotes the full set of fermionic occupation numbers $\{n_{k,a}\}$ specifying a many-body Floquet eigenstate (a Fock basis label) in a fixed parity sector. In particular, occupying both modes in the same k block ($n_{k,+} = n_{k,-} = 1$) gives zero net contribution, as expected from the $\pm$ pairing.

At $\tau = 1$, the reduced block simplifies to

$$\hat{u}_k(1) = e^{-i\frac{\pi}{2}\hat{\sigma}^x} e^{-i\frac{\pi}{2}(\cos k \hat{\sigma}^x + \sin k \hat{\sigma}^y)} = -\left(\cos k \hat{I} + i \sin k \hat{\sigma}^z\right) = -e^{+ik\hat{\sigma}^z}. \quad (5.93)$$

The principal magnitude gives $\varphi(k; 1) = \arccos(-\cos k) = |\pi - k|$ (for $k \in [0, 2\pi)$), which is folded. Our sign rule [Eq. \(5.89\)](#) removes this loss of information analytically: since $s(k) = \text{sgn}(\cos(k/2)) = \text{sgn}(\pi - k)$ on $k \in [0, 2\pi)$, it yields

$$\vartheta(k; 1) = s(k) |\pi - k| = \pi - k \in (-\pi, \pi], \quad (5.94)$$

i.e. the expected linear (SWAP-like) unfolded phase.

Correspondingly, one convenient unfolded single-particle Floquet generator at the DU point is

$$\hat{h}_k^{\text{DU}} = \pi \hat{I} - k \hat{\sigma}^z, \quad \text{so that} \quad \hat{u}_k(1) = e^{-i\hat{h}_k^{\text{DU}}}. \quad (5.95)$$

In second-quantised form this gives a quadratic unfolded DU Floquet Hamiltonian

$$\hat{H}_F^{\text{DU}} = \pi \hat{N}_f + \sum_{k \in K_{\hat{P}}} k (\hat{n}_{k,B} - \hat{n}_{k,A}), \quad \hat{N}_f = \sum_k (\hat{n}_{k,A} + \hat{n}_{k,B}). \quad (5.96)$$

5.6.2 Comparison of unfolding methods




Figure 5.1: Comparison of unfolding methods for the $N = 12$ XX brickwork drive along the τ -adiabatic path. (a) Folded many-body quasienergies ϵ vs. τ and analytically unfolded (dotted lines). (b) Many-body unfolded quasienergies (obtained with MEO) $\tilde{\epsilon}$ vs. τ . The unfolded lines match the analytical prediction. The grey dashed lines mark the margins of the logarithm branches.

This subsection benchmarks the MEO algorithm against two natural references for the brickwork-driven XX chain. We compare three representations of the generator along the τ -adiabatic path: (i) the folded principal-branch construction, (ii) the unfolded construction obtained by the MEO protocol of [Sec. 5.5](#), and (iii) the analytically unfolded construction [Eq. \(5.92\)](#). We compute the many-body quasienergies in [Fig. 5.1](#) of them three and then, verified that the two unfolded methods lead to the same Floquet Hamiltonian, the Pauli profile (locality as in [Eq. \(5.75\)](#) and number of Pauli strings) for the folded and MEO-unfolded, as in [Fig. 5.2](#).

Throughout this comparison we use the ST discretisation of the path with $\tau_{\min} = 0.01$ and $M_\tau = 100$ in [Eq. \(5.67\)](#). That is, we start from a small- τ baseline (which can easily be connected to the continuous time spectrum) and track the sequence of Floquet unitaries $\{\hat{U}(\tau_j)\}$ up to the discrete-time endpoint $\tau = 1$. The MEO procedure produces a continuous (unfolded) eigenphase assignment by matching eigenspaces across steps and resolving near-degeneracy ambiguities by minimising $\text{Loc}(\hat{H}_F)$, yielding a generator $\hat{H}_F(\tau)$ constrained to change as little as possible from step to step along the adiabatic path.

For the data shown below we take $N = 12$, for which the free-fermion reduction has $L = N/2 = 6$ reduced-momentum points per fermion-parity sector. The allowed




Figure 5.2: Comparison of Pauli-string properties for the MEO generator and the 1st logarithm-branch generator for the $N = 12$ XX brickwork drive along the τ -adiabatic path. (a) Average Pauli-string weight, rescaled to the interval $[0, 1]$, $\widetilde{\text{Loc}}(\hat{H}) = (\text{Loc}(\hat{H}) - 2)/(3N/4 - 2)$, versus τ . In (a1) we also show the unshifted quantity $\text{Loc}(\hat{H})$. The principal-branch construction develops strong nonlocality as soon as the spectrum reaches the first logarithmic branch crossing, while MEO yields a substantially more local representation. At the DU endpoint, locality appears to grow sublinearly with system size. (b) Number of nonzero Pauli strings in $\hat{H}$ vs. τ . Across all panels, the principal-branch construction develops strong branch-cut artefacts near $\tau = 1$, while MEO tracks the analytically unfolded reference without using any model-specific integrable structure.

reduced momenta are

$$K_+ = \left\{ \frac{(2m+1)\pi}{L} : m = 0, 1, 2, 3, 4, 5 \right\}, \quad K_- = \left\{ \frac{2m\pi}{L} : m = 0, 1, 2, 3, 4, 5 \right\}, \quad (5.97)$$

corresponding, respectively, to even ($\hat{P} = +1$) and odd ($\hat{P} = -1$) fermion-parity sectors. Each sector contributes 2^{N-1} many-body quasienergies, and the full 2^N spectrum is obtained as the union of the two sectors.

In Fig. 5.1 we show the many-body quasienergies along the τ -path for the three constructions. Lines connect quasienergies associated with eigenprojectors of maximal overlap at consecutive τ values, and colours are inherited from the continuous-time ($\tau=0$) Hamiltonian. In panel (a), the folded (principal-branch) quasienergies ϵ exhibit apparent crossings and discontinuities when eigenphases cross the branch cut, while the analytical unfolding of Sec. 5.6.1 (dotted lines) remains smooth. Panel (b) shows the MEO unfolded spectrum, which matches the analytical unfolding: analytically, branch selection is fixed by the closed-form dispersion, whereas MEO determines the required 2π shifts numerically through eigenspace tracking and locality minimisation. Grey dashed lines mark logarithm-branch boundaries, crossed multiple times along the path. We apply MEO both in symmetry-resolved sectors (magnetisation and translation, reaching $N=12$) and in full-Hilbert-space ED (thick faded lines), obtaining the same unfolding. In Fig. 5.2 we compare Pauli-string properties of folded and unfolded generators. Panel (a) reports the locality functional $\text{Loc}(\hat{H})$, normalised by the continuous-time baseline ($\text{Loc} = 2$) and by the maximal value ($3N/4$), with the unnormalised $\text{Loc}(\hat{H})$ shown in the inset. The folded generator becomes rapidly

nonlocal as $\tau \rightarrow 1$, whereas the MEO generator follows the analytical reference across the path. The normalised locality decreases with increasing N , consistent with a quasi-local DU Hamiltonian from MEO. Panel (b) shows the number of nonzero Pauli strings in $\hat{H}_F$, normalised by 4^N : again, MEO tracks the analytical unfolding, while the folded construction develops branch-cut-induced long-range artefacts.

Overall, this benchmark establishes three main conclusions. First, at the spectral level, MEO recovers the analytically correct unfolding across the full τ -path, including near-degenerate regions, and it agrees with both symmetry-resolved and full-Hilbert-space calculations. Second, at the generator level, principal-branch cuts are the direct source of spurious nonlocality and inflated Pauli support, while the locality-guided MEO selection removes these artefacts and stays quantitatively aligned with the analytic reference. Third, at the DU endpoint, the observed scaling of locality is consistent with a quasi-local unfolded generator in this integrable XX setting. Taken together, these results validate the core MEO principle: among the many branch choices compatible with a given $\hat{U}(\tau)$, locality selects the physically meaningful continuation and yields a reliable Hamiltonian representation for subsequent Krylov analysis.

5.6.3 XX-DU point Krylov complexity




Figure 5.3: Complexity of the XX model MEO derived generator along the adiabatic path, and for the first branch generator at the DU point. The initial operator is $|O_0\rangle = \frac{1}{\sqrt{N}} \sum_n |\sigma_n^x\rangle$.

We observe that the Krylov coefficients $b_m^{(\mathcal{L})}$ terminate after a modest number of steps and that the complexity $K(t)$ enters an oscillatory regime at short times, consistent with the expected finite-dimensional Krylov space. (a) shows the full $b_m^{(\mathcal{L})}$ sequence for $N = 12$ and the DU generator, while (b) shows the corresponding $K(t)$ computed over 500 Lanczos vectors, which saturates at a finite value.

At the dual-unitary endpoint $\tau = 1$, our unfolded Floquet generator $\hat{H}_F^{\text{DU}}$ is still quadratic in Jordan-Wigner fermions (equivalently, bilinear in Majoranas). This has a sharp consequence for Krylov dynamics under the Liouvillian $\mathcal{L} = [\hat{H}_F^{\text{DU}}, \cdot]$.

Let us introduce the Majoranas $\{\hat{\gamma}_\mu\}_{\mu=1}^{2N}$ and the degree- r operator sectors

$$\mathcal{V}_r = \text{span}\{\hat{\gamma}_{\mu_1}\hat{\gamma}_{\mu_2}\cdots\hat{\gamma}_{\mu_r}\}, \quad \dim \mathcal{V}_r = \binom{2N}{r}. \quad (5.98)$$

Because $\hat{H}_F^{\text{DU}}$ is quadratic, $\mathcal{L}$ acts by a linear rotation on Majoranas and therefore preserves the Majorana degree:

$$\mathcal{L}(\mathcal{V}_r) \subseteq \mathcal{V}_r, \quad r = 0, 1, \dots, 2N. \quad (5.99)$$

Hence, for any initial operator $\hat{O}_0 \in \mathcal{V}_r$, the entire Krylov subspace $\mathcal{K}(\hat{O}_0) = \text{span}\{\hat{O}_0, \mathcal{L}\hat{O}_0, \mathcal{L}^2\hat{O}_0, \dots\}$ is finite-dimensional, with

$$D_K(\hat{O}_0) = \dim \mathcal{K}(\hat{O}_0) \leq \binom{2N}{r}. \quad (5.100)$$


In Lanczos form this finiteness appears as an *exact termination* of the recursion: there exists a step n_* such that $b_{n_*} = 0$ and the Krylov chain closes at dimension $D_K = n_*$. Consequently the Krylov amplitudes evolve under a finite tridiagonal matrix and the Krylov complexity $K(t) = \sum_{n \geq 0} n |\phi_n(t)|^2$ is bounded and (for generic incommensurate finite spectra) quasi-periodic, displaying oscillations and recurrences rather than sustained growth.

Numerically, we study the $b_m^{(\mathcal{L})}$ and corresponding $K(t)$ for up to $N = 12$ and the MEO-unfolded $\hat{H}_F^{\text{DU}}$, we observe that the $|O_0\rangle = \frac{1}{\sqrt{N}} \sum_n |\sigma_n^x\rangle$ generates small Krylov orbits, with $b_m^{(\mathcal{L})}$ terminating after a modest number of steps and $K(t)$ entering an oscillatory regime at short times, see Fig. 5.3. This behaviour is *nontrivial*: it shows that even when the unfolded generator is highly nonlocal in real space at the DU point, its *quadratic (free-fermion) structure* strongly constrains operator growth in Krylov space. In particular, the absence of prolonged Krylov growth here should be understood as a structural consequence of the free-fermion DU endpoint rather than as evidence for generic non-interacting behaviour.

5.7 Other DU models


We now apply MEO to additional DU-endpoint models: the XXX brickwork model and the self-dual kicked Ising (SDKI) at $h = \pi/4$ and $h = 1$. In these interacting settings we do not have an exact single-particle reduction, and therefore no closed-form unfolding reference. The XXX model is integrable, whereas the SDKI at $h = \pi/4, 1$ leads to nonintegrable models, i.e. chaotic (though not maximally chaotic due to a hidden symmetry, as discussed in Ref. [330]). While for the XX and XXX the

Figure 5.4: Locality of the folded and MEO-unfolded generators for the XXX model along the τ -adiabatic path. (a) Average Pauli-string weight, rescaled to the interval $[0, 1]$, $\widetilde{\text{Loc}}(\hat{H}) = (\text{Loc}(\hat{H}) - 2)/(3N/4 - 2)$, versus τ . In (a1) we also show the unshifted quantity $\text{Loc}(\hat{H})$. The folded generator rapidly approaches the $3N/4$ ceiling, while the unfolded generator remains at lower locality and evolves more smoothly. Moreover, the unfolded value at the DU point is approximately N -independent, consistent with a quasi-local generator. (b) Number of nonzero Pauli strings in $\hat{H}$ vs. τ , rescaled by 4^N . The number of nonzero Pauli strings quickly saturates to similar values for folded and unfolded Hamiltonians. At larger N the saturation value is reached at larger τ , meaning that for larger N the intensity of the ST kick needs to be stronger to reach the saturation value.



continuous-time limit $\tau \rightarrow 0$ was natural, as the SDKI gate is not usually defined as the exponentiation of a single Hamiltonian, we had to derive one; here, the branch choice already starts to play a role. We discuss the properties of this gate and two possible choices of the initial seed Hamiltonian in Sec. G, and we adopt the choice in Sec. G.3.

Figure 5.5: Comparison of the locality of the MEO and first-branch generators for the $\text{SDKI}_{(\pi/4)}$ model along the τ -adiabatic path for $N = 6, 8, 10, 12$. The locality at the DU point for both folded and MEO-unfolded Floquet Hamiltonians appears to increase abruptly to the $3N/4$ value at around $\tau \simeq 0.4$ for all N . The right panels show the number of nonzero Pauli strings in the generator normalised by 4^N . For both folded and MEO-unfolded cases, the number of nonzero Pauli strings grows and quickly saturates near 4^N , indicating support over essentially the full operator space. Once normalised by 4^N , the saturation value is reached at larger τ for larger N , meaning that a stronger ST kick is needed to reach saturation.



According to the quantum cellular automaton (QCA) viewpoint, brickwork DU circuits built from SWAP-like (or fermionic-SWAP/iSWAP-like) gates implement, on local operators, a lattice translation automorphism (up to on-site phases), and are therefore not expected in general to admit a strictly local effective generator in 1D;

correspondingly, the unfolded generator should develop long-range terms as $\tau \rightarrow 1$ ⁴.




Figure 5.6: Comparison of the locality of the MEO and first-branch generators for the $\text{SDKI}_{(1)}$ model along the τ -adiabatic path for $N = 6, 8, 10, 12$. In panel (a) we show the locality of the folded and MEO-unfolded Floquet Hamiltonians. For the MEO-unfolded case, the locality stays near the initial continuous-time value and then jumps towards $3N/4$, characteristic of random operators. In panel (b) we show the number of nonzero Pauli strings, which grows for both folded and MEO-unfolded Hamiltonians and quickly saturates near 4^N , indicating support over essentially the full operator space. Once normalised by 4^N , the saturation value is reached at larger τ for larger N , meaning that a stronger ST kick is needed to reach saturation.

Taken together, Figs. 5.4 to 5.6 show a clear separation between integrable and chaotic DU families. For the XXX, the MEO-unfolded generator remains comparatively local along the path, and at the DU endpoint its locality grows at most sublinearly with N , consistent with a quasi-local description. For the $\text{SDKI}_{(\pi/4)}$ and $\text{SDKI}_{(1)}$, instead, the locality rapidly approaches the random-operator scale $3N/4$, signaling a strongly nonlocal unfolded generator. Alongside locality, the number of nonzero Pauli strings grows much faster in $\text{SDKI}_{(1)}$ than in the integrable cases, and approaches support over essentially the full operator space near the endpoint.

For the XXX we take as initial seed $|O_0\rangle = \frac{1}{\sqrt{N}} \sum_{n=0}^{N-1} e^{iqn} |\sigma_n^x \sigma_{n+1}^y - \sigma_n^y \sigma_{n+1}^x\rangle$ with $q = 0.1$, which is a local current operator. For the SDKI we take as initial seed $|O_0\rangle = \frac{1}{\sqrt{N}} \sum_{n=0}^{N-1} e^{iqn} |\sigma_n^z\rangle$ with $q = 0.1$. These follow from the ones taken in Ref. [293] for the corresponding continuous-time models.

Overall, MEO yields a continuous unfolding along the full τ -path and keeps the generator local for longer than the first-branch construction. At the DU endpoint, it remains quasi-local for integrable models but develops long-range structure for chaotic ones, in agreement with the QCA-based expectation above.

Finally, we compute Lanczos coefficients and Krylov complexity for $\text{SDKI}_{(\pi/4)}$ and $\text{SDKI}_{(1)}$ along the adiabatic path, including the first-branch Hamiltonian at the DU point (Fig. 5.7). At small τ both models show a similar growth regime, consistent with the fact that they are close to the same continuous-time limit. However, as τ

⁴In 1D, such shift dynamics is a QCA with nontrivial GNVW index, while partitioned finite-depth circuits belong to the trivial-index class [331]. A standard consequence is that translation/shift edge dynamics cannot be generated by purely 1D finite-time Hamiltonian evolution with Lieb-Robinson locality; see the explicit argument in the Floquet-edge context in Ref. [332].




Figure 5.7: Comparison of the Lanczos coefficients and the K-complexity of the MEO generator for the $\text{SDKI}_{(\pi/4)}$ and $\text{SDKI}_{(1)}$ models along the τ -adiabatic path and for the first branch Hamiltonian at the DU point (dashed black line), for $N = 10$ and number of Krylov vectors $M = 500$. The complexity saturates quickly to a finite value for the integrable $\text{SDKI}_{(\pi/4)}$ model while for the chaotic $\text{SDKI}_{(1)}$ it appears to grow similarly to the $\tau = 0$ Hamiltonian for longer.

increases, the $\text{SDKI}_{(\pi/4)}$ model saturates rapidly to finite complexity, while the chaotic $\text{SDKI}_{(1)}$ case shows a longer growth regime, close to the $\tau = 0$ behaviour.

However, the OGH has been stated only for local generators, and therefore the behaviour of the complexity at the DU point for the $\text{SDKI}_{(\pi/4)}$ and $\text{SDKI}_{(1)}$ model should be taken with caution, as it may be influenced by finite-size effects and the specific choice of initial operator. A more detailed analysis of the complexity growth in nonlocal settings would be an interesting direction for future research.

5.8 Conclusion

In “Complexity, chaos, and transport”, we show how the quantities discussed in Part I can be interpreted as forms of complexity diagnostics. In that sense, spectral diagnostics, transport properties, and relaxation speed are different manifestations of how dynamical information is organised in Krylov space. The goal of this Chapter instead is to build a controlled route for studying operator complexity in Floquet circuits, in direct dialogue with continuous-time Krylov theory. The continuous-time Lanczos and discrete-time Arnoldi constructions reviewed in Sec. 5.2 are complementary

descriptions of the same reorganisation of dynamics in Krylov space, consistent with the operator growth hypothesis (OGH) framework [293, 307, 308].

A central conceptual step is the spectral-aliasing identity between unitary and Liouvillian spectral functions (Secs. 5.3 and 5.3.2), which makes explicit the many-to-one nature of Floquet folding at the level of moments [195, 309]. This is precisely why branch selection for $\log \hat{U}$ is not a technical detail but a physical choice. We therefore devise a maximum eigenprojector overlap (MEO) method (Sec. 5.5) which provides such a choice by combining eigenprojector tracking with a locality principle along an adiabatic path, yielding a physically motivated unfolding of quasienergies and therefore a physically motivated Floquet Hamiltonian [313].

This gives a concrete pipeline from Floquet circuits to complexity diagnostics:

$$\hat{U}(\tau) \xrightarrow{\text{MEO}} \hat{H}_F(\tau) \xrightarrow{\mathcal{L}=[\hat{H}_F, \cdot]} \{b_m^{(\mathcal{L})}\}. \quad (5.101)$$

In this sense, MEO makes Floquet complexity analysis more analogous to continuous-time models: once an unfolded generator is selected, one can apply the standard Lanczos/Krylov machinery directly to $\hat{H}_F$, rather than relying only on Arnoldi data of the unitary map.

In the XX benchmark, where an analytic unfolding is available (Secs. 5.6, 5.6.1 and 5.6.2), MEO reproduces the expected branch assignment and generator structure. The unfolded generator remains quasi-local at the DU point, and the resulting Krylov dynamics is bounded, consistent with the underlying free-fermion structure.

Applying the same strategy to interacting DU families (Sec. 5.7) reveals a structural split. For XXX, the unfolded Floquet Hamiltonian remains quasi-local, so Lanczos coefficients and Krylov complexity are quantitatively accessible in close analogy with continuous-time local models. For $\text{SDKI}_{(\pi/4)}$ and $\text{SDKI}_{(1)}$, the DU endpoint is consistent with a substantially nonlocal generator, in line with QCA-based expectations [41]. Yet their complexity behaviour differs: $\text{SDKI}_{(\pi/4)}$ saturates rapidly, while $\text{SDKI}_{(1)}$ exhibits a longer growth regime. Thus nonlocality of the unfolded generator, by itself, does not uniquely determine late-time complexity growth.

This distinction is important because OGH is formulated for local generators [293], and there is no general proof that its conclusions extend to genuinely nonlocal Floquet Hamiltonians. Complexity at DU points without quasi-locality should therefore be interpreted cautiously, and finite-size effects and additional structures (e.g. hidden symmetries) may be decisive. More broadly, once a physically motivated unfolded generator is available, one can probe Floquet-circuit properties beyond Krylov complexity within standard effective-Hamiltonian frameworks (Sec. 5.4) [23, 24, 99], opening a systematic route to analytical and physically grounded descriptions even in interacting



settings.

Chapter 6

Summary and outlook

The role of Floquet circuits in quantum computation, as discussed in [Ch. 1](#), is best understood operationally. On a gate-based device the elementary object is not a Hamiltonian but a *protocol*: a calibrated cycle of local operations. Repeating that cycle implements a discrete-time many-body evolution $\hat{U}^t$. From this viewpoint, Floquet systems are not only a convenient theoretical idealisation, they are the *experimentally primary* description of programmable quantum matter.

Discrete time does two things at once. It discretises time and relaxes energy conservation, but it does *not* remove the organising principles that make many-body physics predictive: locality, symmetries (and corresponding conserved quantities) still control the infrared. The central message of this thesis is that, in this setting, the right diagnostics must be Floquet-native: they should speak the language of quasienergy and stroboscopic sampling, and they should be extractable from measurements a quantum processor actually returns. Within that constraint, we can still link three layers that are usually discussed separately: *spectral structure* (what the spectrum/eigenbasis looks like), *transport and relaxation* (which conserved structures bottleneck late times and intermediate times), and *operator complexity* (how quickly simple operators become hard to describe).

A second message, also already implicit in the introduction, is that *transport, chaos, and scrambling are logically distinct questions* ([Sec. 1.1](#)). In Floquet circuits this separation becomes particularly clean: a unitary can look spectrally chaotic while transport ranges from ballistic to diffusive to anomalous depending on constraints, and finite-time prethermal windows can dominate what an experiment sees even when infinite-time diagnostics look featureless. This thesis also focuses on how to make those separations *measurable*.

For readability we summarise the main achievements as a small set of linked tools and the physics they expose.

Floquet circuits generically lack energy conservation, so transport is organised by whatever local conservation laws remain; here that is U(1) magnetisation in an XXZ Floquet circuit. We study two deformations of the same setting: disorder (spatial and with local Z-fields) and a single impurity [[2](#), [5](#)].

For the disordered circuit we first use spectral diagnostics to map regimes from

less chaotic (localised-like) to ergodic, then switch to finite-time, readout-friendly probes: drift and spread (circular variance) extracted from local-quench Z -profiles. These compressed observables remain stable under disorder averaging and periodic boundaries, and they expose a long-lived near-SWAP *swappy* window and its ergodic-swappy crossover (Sec. 3.8).

For the impurity circuit we show that a single local defect can drastically change spectral properties while leaving the current spectral function essentially unchanged. We generalise a lattice current operator and compute its eigenphase-resolved spectral density using a kernel moment method adapted to discrete-time unitaries; in thermalising regimes this quantity plays the role of conductivity. Finally, quasi-conserved-charge diagnostics reveal a defect-controlled prethermal plateau (Sec. 2.8).

Hybrid
workflow
for spec-
tral recon-
struction

On the methods side, we implemented a moment-expansion (KPM) approach to spectral densities as a hybrid quantum-classical workflow and validated it experimentally by extracting an interacting-spin DOS on a trapped-ion processor up to $N = 18$ spins [1]. The point is not only that one can reconstruct a spectrum; it is that *moments* provide a bridge between what gate-based hardware does natively and what continuous-time theory asks for. On near-term devices the most natural moment problem is Floquet: the implemented object is a discrete-time unitary cycle $\hat{U}$, and estimating moments is essentially estimating powers $\hat{U}^t$ from stroboscopic measurements. Continuous-time Hamiltonian spectroscopy, by contrast, requires an approximation layer (trotterisation) in the near term, and only becomes truly native in the fault-tolerant regime where polynomial signal processing (QSVT) allows direct access to controlled functions of $\hat{H}$; in this sense our quantum-classical KPM workflow for a trotterised Hamiltonian model can be understood as an experimentally accessible approximation to the QSVT moment framework, while simultaneously motivating why Floquet circuits are the most native arena for these ideas on present-day processors. This turns “running a circuit” into a controlled way of probing spectral structure, rather than only producing time traces.

Anomalous
relaxation
as resource
depletion

In Ch. 4 we reframed anomalous equilibration, including thermal and symmetry Mpemba effects, in resource-theoretic terms [3]. The conceptual payoff is a clean criterion: monotone crossings are controlled by overlaps with slow, resourceful modes. This unifies different Mpemba phenomena and clarifies why the symmetry Mpemba effect can arise even through a purely classical mechanism.

We then pushed the same logic beyond thermodynamic/symmetry resources to “complexity-like” resources in unitary circuits, including two-step Pontus-Mpemba protocols [4]. The resulting message fits the thesis theme: even in closed unitary dynamics, the way structure decays *locally* can be organised and accelerated by

understanding which slow modes are actually populated.

Part II asks how operator structure grows under circuit evolution. Discrete time brings a kinematic obstruction: quasienergy periodicity causes spectral folding. We formalised this as a spectral-aliasing identity connecting continuous-time (Lanczos) and discrete-time (Arnoldi) recursion at the level of measures and moments [6].

The practical implication is cautionary: access to $\hat{U}$ alone is not enough to define a faithful continuous-time generator because aliasing obscures the information the OGH is formulated in terms of. Reconnecting to continuous-time Krylov theory requires an effective generator $\hat{H}_F = \log \hat{U}$ together with a physically motivated branch choice. We address this with a locality-guided adiabatic unfolding protocol that tracks eigenprojectors along a Floquet path and selects branches that yield quasi-local $\hat{H}_F$ when possible. At present this is not a hardware-native diagnostic: the unfolding step is computationally demanding and, in generic interacting settings, scales poorly (effectively requiring detailed spectral information).

We benchmark dual-unitary endpoint families (free-fermion XX, XXX, and self-dual kicked Ising (SDKI) circuits) and find qualitatively distinct behaviours. In the interacting SDKI family the unfolded $\hat{H}_F$ can become strongly nonlocal near the dual-unitary point, so OGH-style interpretations (which assume a local generator) require extra care. Within the same family we observe both rapid saturation (e.g. at $h = \pi/4$) and a much longer growth regime (e.g. at $h = 1$), depending on the SDKI gate parameter.

Three themes recur across the thesis.

1. **Floquet-native diagnostics matter.** The cleanest objects are those defined directly in discrete time (quasienergies, stroboscopic correlators, cycle-resolved profiles) and those compressions of data that match experimental readout (moments, circular estimators, finite-time profile probes).
2. **Static chaos is not transport, and neither is operator growth.** Spectral universality is a useful ceiling (“saturated complexity”), but it does not determine the bottlenecks that control transport coefficients or intermediate-time structure. Conversely, transport constraints can strongly shape operator growth even when eigenphases look chaotic.
3. **Intermediate-time structure is often the physics.** Prethermal windows and crossovers are not nuisances to be averaged away; in circuits they are often the dominant phenomenology, and they are organised by quasi-conserved Floquet structures tied to solvable stroboscopic points and to how they are broken (globally vs locally, weakly vs strongly).

6.1 Outlook

The work here suggests several natural directions, both near-term and longer-term.

Scaling hybrid spectral reconstruction

Moment-based spectral reconstruction is a natural interface between theory and hardware because it produces quantities that are (i) symmetry-resolved, (ii) comparable across platforms, and (iii) sensitive to implementation details [1]. For more complicated systems, a key scaling lever is to exploit internal structure explicitly: use exact symmetries and constraints to work in charge sectors (or symmetry blocks) and to construct symmetry-resolved traces/correlators, which reduces both the effective Hilbert-space size and estimator variance [245, 333]. This is particularly natural for digitally engineered models with built-in gauge or non-Abelian structure, where the physically relevant spectrum already lives in a constrained subspace and sector-resolved spectral densities become the right benchmarking object [334]. More ambitiously, operator-resolved and symmetry-resolved spectral measures can serve as a compact *fingerprint* of $\hat{U}_{\text{phys}}$ itself, not only of an intended target model.

Designing “digital matter”

The swappy regime and the defect-controlled plateau point towards a more systematic design philosophy: use solvable stroboscopic points (near-SWAP, dual-unitary endpoints, integrable bulks) as anchors, and engineer controlled departures (disorder, defects, spatial inhomogeneity) to sculpt transport and relaxation. This programme naturally connects to stabilising and diagnosing intrinsically nonequilibrium phases in circuits, including subharmonic response and Floquet time-crystalline signatures, within the same experimentally native language.

OGH for nonlocal Floquet generators

The OGH [293] is formulated for local Hamiltonian evolution, where locality structures operator growth and underpins the link between Krylov dynamics and chaos. In Floquet circuits the experimentally primary object is instead a local unitary cycle $\hat{U}$, and connecting to continuous-time heuristics requires choosing an effective generator $\hat{H}_F = \log \hat{U}$. This choice is physical (and generally non-unique): branch cuts and quasienergy folding can produce $\hat{H}_F$ that is strongly nonlocal even when $\hat{U}$ itself is finite-depth and local [6, 23, 320]. Dual-unitary endpoint families provide a sharp testbed for this tension, since unfolding can yield highly nonlocal generators near special points despite the circuit remaining local [6, 335, 336]. A concrete outcome would be an OGH-style extension that remains predictive for genuinely nonlocal Floquet generators, either by formulating growth criteria directly in discrete time or by combining spectral unwinding with locality diagnostics [6, 325].

Local resource dissipation

The resource-mode picture of Mpemba physics suggests a sharper question for closed many-body dynamics: how do resources dissipate when the system is probed *locally*?

For thermal, symmetry, and complexity-like resources, one should track not only whether a resource decreases globally, but also whether local observables are driven towards the corresponding free manifold while the resource is redistributed into nonlocal degrees of freedom [77, 134]. In the case of magic, this perspective may open a new route to understanding how states are steered towards their closest stabiliser description [288, 337, 338] and how nonstabiliser content is lost locally while persisting in delocalised correlations [77, 134, 266], in close analogy with the local decay of coherence and free-energy signatures during thermalisation [212, 213, 339]. This could also sharpen practical magic estimation beyond small- ε neighbourhoods of the stabiliser polytope, where stabiliser-fidelity based characterisations are currently the most tractable [338].

Building on the previous point, it is natural to ask whether resources admit an internal *mode decomposition*, analogous to the harmonic-analysis “modes of asymmetry” that organise symmetric dynamics [204]. For magic and for (fermionic) non-Gaussianity, the goal would be to go beyond a single scalar monotone and define mode-resolved diagnostics that track where the resource lives, how it is redistributed into delocalised correlations, and how it spreads under circuit evolution [266, 269, 340]. Such a framework would make “resource control” more concrete: one could aim to selectively depopulate the slowest or most resourceful modes (in the spirit of the Mpemba mode picture) and connect to recent attempts to relate nonstabilizerness to an asymmetry viewpoint [341].

Resource
modes

6.2 Closing perspective

In the introduction we argued that Floquet dynamics is becoming a baseline language for the second quantum revolution, in the same sense that simple solvable models become baselines in equilibrium physics: not because they are the final story, but because they are tractable, benchmarkable, and experimentally real. This thesis has treated discrete-time many-body evolution as a first-class object and developed probes that stay meaningful when the fundamental description is stroboscopic. For spectra and transport, several of these probes are explicitly aligned with what present-day hardware returns (moments, stroboscopic correlators, compressed profile diagnostics). For operator complexity, the message is that discrete-time aliasing means that $\hat{U}$ alone in the Arnoldi method does not clear heuristics and to connect to the OGH one needs to construct a physically motivated $\log \hat{U}$. As devices and algorithms improve, the hope is that this combination of hardware-aligned diagnostics and sharp “no-go/obstruction” statements helps turn circuit experiments into quantitative many-body laboratories: diagnosing integrability breaking, prethermalisation, and anomalous relaxation in

driven quantum matter, while keeping the analysis aligned with what the hardware actually does and what it actually measures.

Appendix

A Connection to quantum singular value transformation

In [Sec. 2.9.5](#) we devise a way to implement an approximated version of the Chebyshev moments with very limited resources by reducing them to traces of implementable time evolutions, bypassing the full QSVT machinery. The purpose of this appendix is to place the approximation used in [Sec. 2.9.5](#) within the broader QSVT framework and to clarify what is gained and what is sacrificed by replacing exact signal processing with a spectral-bulk expansion of $\arccos(\hat{H})$ [1].

A.1 QSVT in brief

The input model for QSVT is a block encoding of an operator $\hat{A}$ with $\|\hat{A}\| \leq 1$. A unitary $\hat{U}$ acting on an ancilla register a and a system register s is an (α, a) block encoding of $\hat{A}$ if, for some $\alpha \geq 1$,

$$(\langle 0|_{(a)} \otimes \hat{I}_{(s)}) \hat{U} (|0\rangle_{(a)} \otimes \hat{I}_{(s)}) = \frac{\hat{A}}{\alpha}. \quad (1)$$

Given such an embedding, quantum signal processing (QSP) and its generalisation QSVT construct, from alternating uses of $\hat{U}$ and simple single-qubit rotations on the ancilla, a new unitary whose top-left block is (approximately) a desired polynomial $p(\hat{A})$, with query complexity proportional to $\deg(p)$ and near-optimal dependence on the target precision [148–150, 342]. For Hermitian $\hat{A}$, the relevant constraint is that $p(x)$ is bounded on $x \in [-1, 1]$ (and satisfies a parity constraint tied to $\deg(p)$), ensuring the existence of a valid QSP phase sequence [149, 150].

In this language, qubitisation is a particular QSVT instantiation tailored to Hamiltonian simulation. Starting from a linear combination of unitaries (LCU) decomposition $\hat{H} = \sum_p \omega_p \hat{H}_p$ and the corresponding *prepare/select* oracles, one builds a block encoding of the normalised operator $\hat{h} = \hat{H}/\alpha$ and a related *signal unitary* (quantum walk) whose eigenphases encode $\arccos(E_k)$ when E_k are eigenvalues of $\hat{h}$ [148, 149].

A.2 Chebyshev polynomials from the signal unitary

Let $\hat{h}$ be an already normalised Hermitian operator so that $\text{spec}(\hat{h}) \subset (-1, 1)$ and eigenpairs $\hat{h}|k\rangle = E_k|k\rangle$. Given a Hermitian block encoding $\hat{U}$ of $\hat{h}$ as in [Eq. \(1\)](#), QSVT associates to it a *signal unitary* $\hat{W}$ (often realised via the qubitisation walk construction in Hamiltonian simulation, also built from *prepare/select* oracles) whose action reduces, for each k , to a two-dimensional rotation with angle $\arccos(E_k)$ [148–

150].

More precisely, define $|\phi_k\rangle = |0\rangle_{(a)} \otimes |k\rangle_{(s)}$. For $|E_k| < 1$ let $|\phi_k^\perp\rangle$ be the normalised component of $\hat{W}|\phi_k\rangle$ orthogonal to $|\phi_k\rangle$, i.e.

$$\hat{W}|\phi_k\rangle = E_k|\phi_k\rangle + \sqrt{1 - E_k^2}|\phi_k^\perp\rangle, \quad \langle\phi_k|\phi_k^\perp\rangle = 0, \quad \|\phi_k^\perp\| = 1, \quad (2)$$

and set $\mathcal{H}_k = \text{span}\{|\phi_k\rangle, |\phi_k^\perp\rangle\}$. Standard analyses of the signal unitary imply that $\mathcal{H}_k$ is $\hat{W}$ -invariant and that, in the basis $\{|\phi_k\rangle, |\phi_k^\perp\rangle\}$, the restriction $\hat{W}|_{\mathcal{H}_k}$ takes the form [148, 149]:

$$\hat{W}|_{\mathcal{H}_k} = \begin{pmatrix} E_k & \sqrt{1 - E_k^2} \\ -\sqrt{1 - E_k^2} & E_k \end{pmatrix} = e^{i\hat{Y}_{(k)} \arccos(E_k)}, \quad \hat{Y}_{(k)} = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}_{\{|\phi_k\rangle, |\phi_k^\perp\rangle\}}. \quad (3)$$

If $|E_k| = 1$, the rotation angle is 0 or π and the relevant invariant space is one-dimensional.

Since $\hat{W}|_{\mathcal{H}_k}$ is a rotation by $\arccos(E_k)$, its m -th power satisfies

$$(\langle 0|_{(a)} \otimes \hat{I}) \hat{W}^m (|0\rangle_{(a)} \otimes \hat{I}) |k\rangle = \cos(m \arccos(E_k)) |k\rangle = T_m(E_k) |k\rangle, \quad (4)$$

and therefore, as an operator identity on the system register,

$$(\langle 0|_{(a)} \otimes \hat{I}) \hat{W}^m (|0\rangle_{(a)} \otimes \hat{I}) = T_m(\hat{h}). \quad (5)$$

A.3 Spectral-bulk arccosine approximation as a NISQ surrogate

The moment estimator used in Sec. 2.9.5 retains the Chebyshev-moment target but replaces the exact eigenphase map $\arccos(\hat{h})$ by a low-order approximation around the spectral bulk. Using the odd-power expansion around 0,

$$\arccos(x) = \frac{\pi}{2} - \sum_{k=0}^{\infty} c_k x^{2k+1}, \quad c_k = \frac{(2k)!}{2^{2k}(k!)^2(2k+1)}, \quad (6)$$

we truncate to obtain $\arccos(\hat{h}) \approx \frac{\pi}{2} - \hat{h}_K$ with $\hat{h}_K = \sum_{k=0}^K c_k \hat{h}^{2k+1}$. This induces a *proxy* for the Chebyshev polynomial,

$$T_m^{(K)}(\hat{h}) := \cos\left(m\left(\frac{\pi}{2} - \hat{h}_K\right)\right), \quad (7)$$

and the corresponding proxy moments $\mu_m^K = \frac{1}{D} \text{Tr}[T_m^{(K)}(\hat{h})]$, with $D = \dim \mathcal{H}$. For integer m , the known phase factors reduce Eq. (7) to real or imaginary parts of an implementable unitary:

$$\cos\left(m\left(\frac{\pi}{2} - \hat{h}_K\right)\right) = \text{Re}\left[e^{im\pi/2} e^{-im\hat{h}_K}\right], \quad \sin\left(m\left(\frac{\pi}{2} - \hat{h}_K\right)\right) = \text{Im}\left[e^{im\pi/2} e^{-im\hat{h}_K}\right]. \quad (8)$$

This is the bridge used in Sec. 2.9.5: rather than implementing a QSP phase sequence that directly block-encodes $T_m(\hat{h})$ (or $T_m^{(K)}(\hat{h})$), we estimate traces of $e^{-im\hat{h}_K}$ with a DQC1-like circuit.

Two remarks clarify the relationship to QSVT. (i) Expanding the cosine in Eq. (7) shows that $T_m^{(K)}(x)$ is a bounded real polynomial on $[-1, 1]$ (of degree $m(2K+1)$) whenever $\hat{h}_K$ is a polynomial in $\hat{h}$. Hence it lies in the class of transforms implementable by QSVT using $O(m(2K+1))$ queries to the block encoding [149, 150]. Our approach can therefore be viewed as replacing the QSVT polynomial-implementation step by direct Hamiltonian simulation of $e^{-im\hat{h}_K}$. (ii) The power series Eq. (6) is controlled when $|x| \ll 1$ and deteriorates near $|x| = 1$. This matches the KPM regime of interest in Sec. 2.9.5, where the density of states concentrates near the spectral centre after normalisation. In contrast, QSVT is typically paired with near-minimax Chebyshev-series approximations that provide uniform control over $x \in [-1, 1]$ at the cost of higher polynomial degree and deeper phase sequences [149, 150]. A simple bulk error estimate follows from Lipschitz continuity of $\cos$: for any $x \in [-1, 1]$,

$$|T_m(x) - T_m^{(K)}(x)| = |\cos(m \arccos x) - \cos(m(\frac{\pi}{2} - p_K(x)))| \leq m |\arccos x - (\frac{\pi}{2} - p_K(x))|, \quad (9)$$

with $p_K(x) = \sum_{k=0}^K c_k x^{2k+1}$. Thus, in the regime $|x| \leq r < 1$ the proxy-polynomial error inherits the truncation error of the series for $\arccos(x)$ and grows at most linearly with m for fixed K .

A.4 What is preserved and what is traded away

The proxy-moment method of Sec. 2.9.5 preserves the central QSVT insight that Chebyshev moments are natural spectral probes, but it trades the algorithmic optimality of QSVT for a NISQ-feasible primitive:

- In fault-tolerant settings, QSVT implements $T_m(\hat{h})$ (and far more general $p(\hat{h})$) with query complexity $O(m)$ and polylogarithmic dependence on the target precision, using a constant-size ancilla [148–150, 342].
- Here, we avoid long QSP phase sequences and the associated coherent control overhead (which is challenging on current hardware [343]) by estimating $\text{Tr}[e^{-im\hat{h}_K}]$

via a DQC1-like Hadamard test and compiling $e^{-im\hat{h}_K}$ through product-formula simulation. The price is that the resource scaling is set by Hamiltonian-simulation depth (and by shot noise), rather than by the asymptotically optimal QSVT query complexity.

In summary, the arccosine truncation used in [Sec. 2.9.5](#) can be interpreted as a constrained, bulk-optimised surrogate for the QSVT eigenphase map $\arccos(\hat{h})$. Conceptually, the method remains within the QSVT paradigm (polynomial spectral transforms), but it replaces QSVT’s polynomial-implementation step by direct simulation of an implementable unitary whose real and imaginary parts reproduce the desired proxy moments.

B U(1) two-qubit gate parametrisations

Throughout this appendix we use the local spin-1/2 operators $\hat{s}_n^\alpha$ ($\alpha = x, y, z$) with $\hat{s}_n^\pm = \hat{s}_n^x \pm i\hat{s}_n^y$, and (when needed) Pauli operators $\hat{\sigma}_n^\alpha = 2\hat{s}_n^\alpha$. We fix the two-qubit computational basis as $\{|\uparrow\uparrow\rangle, |\uparrow\downarrow\rangle, |\downarrow\uparrow\rangle, |\downarrow\downarrow\rangle\}$.

B.1 From four local Z -phases to (h_n, h'_n, ϕ_n)

A convenient “digital” factorisation of a generic U(1)-symmetric gate is

$$\hat{v}_{n,n+1} = e^{-i\theta_1^{(n)}\hat{s}_n^z} e^{-i\theta_2^{(n)}\hat{s}_{n+1}^z} e^{-i\frac{J}{2}(\hat{s}_n^+\hat{s}_{n+1}^- + \hat{s}_n^-\hat{s}_{n+1}^+)} e^{-iJ_z\hat{s}_n^z\hat{s}_{n+1}^z} e^{-i\theta_3^{(n)}\hat{s}_n^z} e^{-i\theta_4^{(n)}\hat{s}_{n+1}^z}. \quad (10)$$

Since $\hat{s}_n^z\hat{s}_{n+1}^z$ commutes with $\hat{s}_n^z$ and $\hat{s}_{n+1}^z$, the only nontrivial rearrangement concerns the exchange term. Using

$$e^{i(a\hat{s}_n^z + b\hat{s}_{n+1}^z)} \hat{s}_n^\pm \hat{s}_{n+1}^\mp e^{-i(a\hat{s}_n^z + b\hat{s}_{n+1}^z)} = e^{\pm i(a-b)} \hat{s}_n^\pm \hat{s}_{n+1}^\mp, \quad (11)$$

one can “push” the first pair of Z -rotations through the exchange to generate a Peierls phase. Writing

$$\phi_n = \theta_1^{(n)} - \theta_2^{(n)}, \quad h_n = \theta_1^{(n)} + \theta_3^{(n)}, \quad h'_n = \theta_2^{(n)} + \theta_4^{(n)}, \quad (12)$$

[Eq. \(10\)](#) is equivalent (up to regrouping commuting Z -rotations) to the parametrisation used in [Eq. \(2.54\)](#) of the main text,

$$\hat{u}_{n,n+1} = e^{-i\hat{h}_{n,n+1}} e^{-ih_n\hat{s}_n^z} e^{-ih'_n\hat{s}_{n+1}^z}, \quad \hat{h}_{n,n+1} = \frac{J}{2}(\hat{s}_n^+\hat{s}_{n+1}^- e^{i\phi_n} + \hat{s}_n^-\hat{s}_{n+1}^+ e^{-i\phi_n}) + J_z\hat{s}_n^z\hat{s}_{n+1}^z. \quad (13)$$

All angles are understood modulo 2π ; if the $\theta_k^{(n)}$ are sampled uniformly in $[-\pi, \pi)$ and then reduced modulo 2π , the induced (h_n, h'_n, ϕ_n) are also uniform modulo 2π .

B.2 Explicit 4×4 matrix of $\hat{u}_{n,n+1}$

In the basis $\{|\uparrow\uparrow\rangle, |\uparrow\downarrow\rangle, |\downarrow\uparrow\rangle, |\downarrow\downarrow\rangle\}$, the gate $\hat{u}_{n,n+1}$ of Eq. (13) is block diagonal, with a nontrivial 2×2 block on the one-excitation subspace. Writing $c = \cos \frac{J}{2}$, $s = \sin \frac{J}{2}$, and $h_{\pm} = \frac{h_n \pm h'_n}{2}$, one finds

$$\hat{u}_{n,n+1} = \begin{pmatrix} e^{-i\left(\frac{J_z}{4} + h_+\right)} & 0 & 0 & 0 \\ 0 & c e^{+i\left(\frac{J_z}{4} - h_-\right)} & -i s e^{+i\left(\frac{J_z}{4} - \phi_n + h_-\right)} & 0 \\ 0 & -i s e^{+i\left(\frac{J_z}{4} + \phi_n - h_-\right)} & c e^{+i\left(\frac{J_z}{4} + h_-\right)} & 0 \\ 0 & 0 & 0 & e^{-i\left(\frac{J_z}{4} - h_+\right)} \end{pmatrix}. \quad (14)$$

This form makes U(1) symmetry manifest: the corners (zero- and two-excitation sectors) are phases, and the one-excitation sector is a 2×2 unitary.

B.3 Mapping to the Žnidarič-Duh-Zadnik (ZDZ) parametrisations

Žnidarič-Duh-Zadnik (ZDZ) consider homogeneous U(1) two-qubit gates and show that (after removing phases associated with conserved operators) they can be expressed in a Hurwitz (Haar) parametrisation and then mapped to a unitary asymmetric six-vertex $\hat{R}$ -matrix solving the Yang-Baxter equation. We summarise here the identifications needed to translate our parameters $(J, J_z, h_n, h'_n, \phi_n)$ into their conventions.

(i) Gauge-fixing the conserved magnetisation phase. The operator $\hat{s}_n^z + \hat{s}_{n+1}^z$ is conserved by any U(1) gate, so one may right-multiply $\hat{u}_{n,n+1}$ by a local rotation $e^{+i\mu(\hat{s}_n^z + \hat{s}_{n+1}^z)}$ without changing any magnetisation-sector dynamics (it only fixes a parametrisation convention). A convenient choice is $\mu = \frac{h_n + h'_n}{2}$, which equalises the two corner phases in Eq. (14). Define the gauge-fixed gate

$$\hat{u}_{n,n+1}^{(0)} = \hat{u}_{n,n+1} e^{+i\frac{h_n + h'_n}{2}(\hat{s}_n^z + \hat{s}_{n+1}^z)}. \quad (15)$$

Then $\hat{u}_{n,n+1}^{(0)}$ has equal corner phases $e^{-iJ_z/4}$.

(ii) Match to the Hurwitz U(1) form. ZDZ write a generic U(1) two-qubit unitary (with equal corners) as

$$\hat{u}_H = \begin{pmatrix} e^{i\delta} & 0 & 0 & 0 \\ 0 & e^{i\alpha} \sin \varphi_H e^{-i\chi} & e^{i\alpha} \cos \varphi_H e^{-i\vartheta} & 0 \\ 0 & e^{i\alpha} \cos \varphi_H e^{+i\vartheta} & -e^{i\alpha} \sin \varphi_H e^{+i\chi} & 0 \\ 0 & 0 & 0 & e^{i\delta} \end{pmatrix}. \quad (16)$$

Comparing Eqs. (14) and (15) with Eq. (16) yields one consistent identification (modulo 2π)

$$\varphi_H = \frac{\pi}{2} - \frac{J}{2}, \quad \delta = -\frac{J_z}{4}, \quad \alpha = \frac{J_z}{4} + \frac{\pi}{2}, \quad \chi = \frac{h_n - h'_n}{2} + \frac{\pi}{2}, \quad \vartheta = \phi_n - \frac{h_n - h'_n}{2} + \pi. \quad (17)$$

Different (but equivalent) choices are possible depending on basis conventions in the one-excitation subspace; the choice Eq. (17) is adapted to the basis stated at the beginning of this appendix.

(iii) Map Hurwitz angles to the asymmetric six-vertex $\hat{R}$ -matrix. ZDZ introduce a unitary $\hat{R}$ -matrix of the form

$$\hat{R}(u, \rho, \beta, \xi, \theta) = e^{i\beta u} \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & i b(u, \rho) e^{-i\xi u} & -a(u, \rho) e^{-i\theta} & 0 \\ 0 & -a(u, \rho) e^{+i\theta} & i b(u, \rho) e^{+i\xi u} & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \quad (18)$$

with two analytic branches (“phases”) for the weights a, b :

$$\text{phase I: } a(u, \rho) = \frac{\sin u}{\sin(u + i\rho)}, \quad b(u, \rho) = \frac{\sinh \rho}{\sin(u + i\rho)}, \quad (19)$$

$$\text{phase II: } a(u, \rho) = \frac{\sinh u}{\sinh(u + i\rho)}, \quad b(u, \rho) = \frac{\sin \rho}{\sinh(u + i\rho)}. \quad (20)$$

In the conventions of Eq. (16), one may take

$$\theta = \vartheta, \quad \gamma = \delta - \alpha + \pi, \quad (21)$$

and then solve for real (u, ρ) on the branch (phase I or II) that yields a real solution (the separation of the two branches is the critical manifold in ZDZ). Once (u, ρ) are

fixed, the remaining parameters follow from the phases of the diagonal one-excitation entries,

$$\xi = \frac{\chi - \pi/2}{u}, \quad \beta = \frac{\delta}{u}, \quad (22)$$

(with the understood limiting procedure if $u \rightarrow 0$). In our numerics we implement this inversion at the level of the 2×2 one-excitation block, choosing the branch that makes u and ρ real, and then extracting (β, ξ) from Eq. (22). This provides a direct dictionary between the five-parameter gate family $(J, J_z, h_n, h'_n, \phi_n)$ and the ZDZ $\hat{R}$ -matrix parameters.

B.4 Simulation-friendly decompositions

For hardware-oriented simulation it is useful to isolate the *fixed* two-qubit entangler (set by (J, J_z)) from the *gate-dependent* Z -rotations (set by (h_n, h'_n, ϕ_n)). Two observations enable this.

First, rewrite the interaction generator using Pauli operators:

$$\hat{h}_{n,n+1} = \frac{J}{4} \left[\cos \phi_n (\hat{\sigma}_n^x \hat{\sigma}_{n+1}^x + \hat{\sigma}_n^y \hat{\sigma}_{n+1}^y) + \sin \phi_n (\hat{\sigma}_n^x \hat{\sigma}_{n+1}^y - \hat{\sigma}_n^y \hat{\sigma}_{n+1}^x) \right] + \frac{J_z}{4} \hat{\sigma}_n^z \hat{\sigma}_{n+1}^z. \quad (23)$$

Second, note that the Peierls phase can be generated by conjugation with *single-qubit* Z -rotations:

$$e^{-i\frac{\phi_n}{4}(\hat{\sigma}_n^z - \hat{\sigma}_{n+1}^z)} (\hat{\sigma}_n^x \hat{\sigma}_{n+1}^x + \hat{\sigma}_n^y \hat{\sigma}_{n+1}^y) e^{+i\frac{\phi_n}{4}(\hat{\sigma}_n^z - \hat{\sigma}_{n+1}^z)} = \cos \phi_n (\hat{\sigma}_n^x \hat{\sigma}_{n+1}^x + \hat{\sigma}_n^y \hat{\sigma}_{n+1}^y) + \sin \phi_n (\hat{\sigma}_n^x \hat{\sigma}_{n+1}^y - \hat{\sigma}_n^y \hat{\sigma}_{n+1}^x). \quad (24)$$

Therefore, defining the *phase-free* entangler

$$\hat{w}(J, J_z) = \exp \left[-i\frac{J}{4} (\hat{\sigma}_n^x \hat{\sigma}_{n+1}^x + \hat{\sigma}_n^y \hat{\sigma}_{n+1}^y) - i\frac{J_z}{4} \hat{\sigma}_n^z \hat{\sigma}_{n+1}^z \right], \quad (25)$$

the full gate can be implemented as

$$\begin{aligned} \hat{u}_{n,n+1} &= \hat{w}_\phi(J, J_z; \phi_n) e^{-i\frac{h_n}{2}\hat{\sigma}_n^z} e^{-i\frac{h'_n}{2}\hat{\sigma}_{n+1}^z}, \\ \hat{w}_\phi(J, J_z; \phi_n) &= e^{-i\frac{\phi_n}{4}(\hat{\sigma}_n^z - \hat{\sigma}_{n+1}^z)} \hat{w}(J, J_z) e^{+i\frac{\phi_n}{4}(\hat{\sigma}_n^z - \hat{\sigma}_{n+1}^z)}. \end{aligned} \quad (26)$$

Because $\hat{\sigma}_n^x \hat{\sigma}_{n+1}^x$, $\hat{\sigma}_n^y \hat{\sigma}_{n+1}^y$, and $\hat{\sigma}_n^z \hat{\sigma}_{n+1}^z$ mutually commute, $\hat{w}(J, J_z)$ can be precomputed once per (J, J_z) (or compiled once to a native two-qubit primitive), while the disorder enters only through the single-qubit Z -rotations in Eq. (26).

C Transport protocol technicalities

This appendix connects the magnetisation profile protocol of [Sec. 3.9.2](#) to standard dynamical correlation functions and stochastic trace estimation. The aim is to justify interpreting $p_n(t)$ as a normalised version of an equilibrium correlator in the chosen symmetry sector.

C.1 Dynamical correlation function in the full Hilbert space

Transport coefficients and scaling exponents are conventionally formulated in terms of spatiotemporal correlators. For systems conserving $\hat{S}^z$ one considers $C_{n,n'}(t) = \text{Tr}[\hat{s}_n^z(t) \hat{s}_{n'}^z]$. Using $\hat{s}_{n'}^z = \hat{P}_{n'}^\dagger - \hat{I}/2$, one finds $C_{n,n'}(t) = \text{Tr}[\hat{s}_n^z(t) \hat{P}_{n'}^\dagger] - \frac{1}{2} \text{Tr}[\hat{s}_n^z(t)]$. The second term vanishes since $\text{Tr}[\hat{s}_n^z(t)] = 0$, hence

$$C_{n,n'}(t) = \text{Tr}[\hat{s}_n^z(t) \hat{P}_{n'}^\dagger]. \quad (27)$$

Explicit traces are expensive at large Hilbert-space dimension, so one uses stochastic trace estimation, see [Sec. 2.9.3](#). For random vectors $|\psi_r\rangle$ drawn from a unitary-invariant ensemble,

$$\text{Tr}[\hat{A}] \simeq \frac{1}{R} \sum_{r=1}^R \langle \psi_r | \hat{A} | \psi_r \rangle, \quad \text{error} \sim \mathcal{O}\left(\frac{1}{\sqrt{RD}}\right), \quad (28)$$

where D is the total Hilbert-space dimension. Applying this to [Eq. \(27\)](#) and moving to the Schrödinger picture gives $C_{n,n'}(t) \simeq \langle \psi_{\text{rand}} | \hat{P}_{n'}^\dagger (\hat{U}^\dagger)^t \hat{s}_n^z \hat{U}^t \hat{P}_{n'}^\dagger | \psi_{\text{rand}} \rangle$, which can be interpreted as measuring $\hat{s}_n^z$ along trajectories generated from a projected typical state.

C.2 Restriction to the fixed-M subspace

In this thesis we restrict to $\mathcal{H}^{(M)}$. Define the sector-resolved correlator as

$$C_{n,n'}^{(M)}(t) = \text{Tr}[\hat{P}_M \hat{s}_n^z(t) \hat{s}_{n'}^z]. \quad (29)$$

Since $\text{Tr}[\bullet] = \sum_M \text{Tr}[\hat{P}_M \bullet]$, the full-space correlator decomposes as

$$C_{n,n'}(t) = \sum_{M=-N/2}^{N/2} C_{n,n'}^{(M)}(t). \quad (30)$$

Using $\hat{s}_{n'}^z = \hat{P}_{n'}^\dagger - \hat{I}/2$ again yields $C_{n,n'}^{(M)}(t) = \text{Tr}[\hat{P}_M \hat{s}_n^z(t) \hat{P}_{n'}^\dagger] - \frac{1}{2} \text{Tr}[\hat{P}_M \hat{s}_n^z(t)]$. The second trace is fixed by symmetry. Since $[\hat{P}_M, \hat{U}] = 0$ and $\hat{P}_M$ is invariant

under site permutations, $\text{Tr}[\hat{P}_M \hat{s}_n^z(t)] = \text{Tr}[\hat{P}_M \hat{s}_n^z] = \frac{1}{N} \text{Tr}[\hat{P}_M \hat{S}^z] = \frac{M D_N^{(M)}}{N}$ with $D_N^{(M)} = \dim \mathcal{H}^{(M)}$. Therefore

$$C_{n,n'}^{(M)}(t) = \text{Tr}[\hat{P}_M \hat{s}_n^z(t) \hat{P}_{n'}^\dagger] - \frac{M D_N^{(M)}}{2N}. \quad (31)$$

Introducing the mixed initial operator

$$\hat{\rho}_M(0) = \hat{P}_M \hat{P}_{n'}^\dagger \hat{P}_M \quad (32)$$

(which coincides with Eq. (3.4) up to normalisation and the choice $n' = n_0$), we can rewrite Eq. (31) as

$$C_{n,n'}^{(M)}(t) = \text{Tr}[\hat{s}_n^z(t) \hat{\rho}_M(0)] - \frac{M D_N^{(M)}}{2N}. \quad (33)$$

C.3 Typicality within $\mathcal{H}^{(M)}$ and the $M = 0$ semiprobability

Let $|\psi_{\text{rand}}^{(M)}\rangle$ be a typical random state in $\mathcal{H}^{(M)}$. Then, for operators supported on $\mathcal{H}^{(M)}$,

$$\langle \psi_{\text{rand}}^{(M)} | \hat{A} | \psi_{\text{rand}}^{(M)} \rangle \simeq \frac{1}{D_N^{(M)}} \text{Tr}[\hat{P}_M \hat{A}]. \quad (34)$$

Applying this to the first term in Eq. (33) yields the practical estimator

$$\text{Tr}[\hat{s}_n^z(t) \hat{\rho}_M(0)] \simeq \frac{D_N^{(M)}}{\mathcal{N}_M} \langle \psi_{\text{rand}}^{(M)} | \hat{P}_{n'}^\dagger \hat{U}^{-t} \hat{s}_n^z \hat{U}^t \hat{P}_{n'} | \psi_{\text{rand}}^{(M)} \rangle, \quad (35)$$

where $\mathcal{N}_M = \text{Tr}[\hat{P}_M \hat{P}_{n'}^\dagger]$ accounts for the projection normalisation.

In the $M = 0$ sector, the additive constant in Eq. (33) vanishes. This is why the $M = 0$ magnetisation profile $M_n(t)$ is naturally interpreted as a sector-resolved correlator with a local-quench initial condition, and why the affine map Eq. (3.6) produces a unit-mass semiprobability profile suitable for moment-based transport diagnostics.

D Spectral properties and light-cone correlations from a transfer matrix

Along the DU line, dynamical correlations simplify due to a transfer-matrix structure. For light-cone observables one may express correlators

$$C_{n,n'}(t) = \text{Tr}[\hat{s}_n^z(t) \hat{s}_{n'}^z] \quad (36)$$




Figure 1: Properties of a single U(1) gate extracted from transfer-matrix eigenvalues. Panel (a) shows λ_z versus J (population mixing), while panel (b) shows $|\lambda_{\pm}|$ increasing towards the SWAP point.

in terms of the action of a single-site completely positive map (transfer matrix)

$$\mathcal{M}[\hat{\rho}] = \text{Tr}_1 \left[\hat{U}_{0,1} (\hat{\rho} \otimes \hat{I}_2) \hat{U}_{0,1}^\dagger \right], \quad (37)$$

which encodes the effective evolution of one local degree of freedom under the two-site gate. In DU circuits, light-cone correlations are the fundamental non-trivial correlators that survive maximal entanglement; see Ref. [41] and references therein.

Diagonalising $\mathcal{M}$ yields four eigenvalues λ_a and eigenmodes $\hat{r}_a$. For the family considered here, the eigenvalues can be written as

$$\lambda_0 = 1, \quad \lambda_z = \frac{1 - \cos J}{2}, \quad \lambda_{\pm} = \pm \sin \frac{J}{2} \sin \frac{J_z}{2}, \quad (38)$$

with corresponding eigenmodes

$$\hat{r}_0 = \hat{I}_2, \quad \hat{r}_z = \hat{s}^z, \quad \hat{r}_{\pm} = \begin{pmatrix} 0 & e^{i(h_1 - \varphi_0)} \\ e^{-i(h_1 - \varphi_0)} & 0 \end{pmatrix}. \quad (39)$$

Crucially, the eigenvalue spectrum is independent of the random phases h_0 , h_1 , and φ_0 , which appear only in the eigenmodes $\hat{r}_{\pm}$. This phase-independence enables exact computations of light-cone correlators up to finite-size effects.

The entire DU family with $J = \pi$ displays non-ergodic dynamics in the sense of Ref. [41], characterised by the presence of at least two independent conserved quantities. Approaching the DU line, the circuit inherits (near-)conserved operators that effectively fragment finite-size spectra into overlapping symmetry sectors, distorting raw level-spacing statistics unless these sectors are resolved [344].

E From global G-invariance to local G-covariance

In this appendix, we demonstrate how a global symmetry of a joint system-environment unitary evolution induces covariance of the reduced system dynamics under the same symmetry group, if the environment is prepared in a symmetric state. Let us consider

a group G , a G -invariant unitary $\hat{T} \in \mathcal{B}(\mathcal{H}_s \otimes \mathcal{H}_e)$, and a G -invariant state $\hat{\pi}_e \in \mathcal{S}(\mathcal{H}_e)$. Then, the map

$$\mathcal{E}[\cdot] = \text{Tr}_e[\hat{T}([\cdot] \otimes \hat{\pi}_e)\hat{T}^\dagger] \quad (40)$$

will be G -covariant, namely

$$\mathcal{E}[\hat{U}_{g,s} \hat{\rho} \hat{U}_{g,s}^\dagger] = \hat{U}_{g,s} \mathcal{E}[\hat{\rho}] \hat{U}_{g,s}^\dagger \quad \forall \hat{\rho} \in \mathcal{S}(\mathcal{H}_s), \quad (41)$$

where $\hat{U}_{g,s}$ is the unitary representation of G on $\mathcal{H}_s$.

To prove this let us consider a unitary representation of the group $\hat{U}_g = \hat{U}_{g,s} \otimes \hat{U}_{g,e}$ on $\mathcal{H}_s \otimes \mathcal{H}_e$ and begin from the left-hand side

$$\mathcal{E}[\hat{U}_{g,s} \hat{\rho} \hat{U}_{g,s}^\dagger] = \text{Tr}_e[\hat{T}(\hat{U}_{g,s} \hat{\rho} \hat{U}_{g,s}^\dagger \otimes \hat{\pi}_e)\hat{T}^\dagger]. \quad (42)$$

By G -invariance we have $[\hat{T}, \hat{U}_{g,s} \otimes \hat{U}_{g,e}] = 0 \forall g \in G$, so

$$\hat{T}(\hat{U}_{g,s} \hat{\rho} \hat{U}_{g,s}^\dagger \otimes \hat{\pi}_e)\hat{T}^\dagger = \hat{U}_{g,s} \otimes \hat{U}_{g,e} (\hat{\sigma}) \hat{U}_{g,s}^\dagger \otimes \hat{U}_{g,e}^\dagger, \quad (43)$$

where $\hat{\sigma} = \hat{T}(\hat{\rho} \otimes \hat{\pi}_e)\hat{T}^\dagger$, using $\hat{\pi}_e = \hat{U}_{g,e} \hat{\pi}_e \hat{U}_{g,e}^\dagger$. Hence

$$\mathcal{E}[\hat{U}_{g,s} \hat{\rho} \hat{U}_{g,s}^\dagger] = \text{Tr}_e[\hat{U}_{g,s} \otimes \hat{U}_{g,e} (\hat{\sigma}) \hat{U}_{g,s}^\dagger \otimes \hat{U}_{g,e}^\dagger]. \quad (44)$$

Since the partial trace over the environment commutes with unitaries on the environment, one gets

$$\begin{aligned} \text{Tr}_e[\hat{U}_{g,s} \otimes \hat{U}_{g,e} (\hat{\sigma}) \hat{U}_{g,s}^\dagger \otimes \hat{U}_{g,e}^\dagger] &= \hat{U}_{g,s} \text{Tr}_e[\hat{\sigma}] \hat{U}_{g,s}^\dagger \\ &= \hat{U}_{g,s} \mathcal{E}[\hat{\rho}] \hat{U}_{g,s}^\dagger. \end{aligned} \quad (45)$$

Thus

$$\mathcal{E}[\hat{U}_{g,s} \hat{\rho} \hat{U}_{g,s}^\dagger] = \hat{U}_{g,s} \mathcal{E}[\hat{\rho}] \hat{U}_{g,s}^\dagger, \quad (46)$$

showing that $\mathcal{E}$ is G -covariant whenever $\hat{T}$ is G -invariant and $\hat{\pi}_e$ is G -invariant.

F Asymptotic decay of the quantum relative entropy

For both classical and quantum systems, the late time behaviour of a state is determined by its overlap with the slowest decaying eigenmode, among the whole spectrum for athermality and only among the nonzero- μ -sectors for asymmetry. Here we prove the equations Eq. (4.19) and Eq. (4.26).

F.1 Proof of the Asymptotic Decay of the KL-Divergence

Let $\hat{\mathcal{L}}^{\text{cl}}$ be a classical Liouvillian with unique stationary distribution $\boldsymbol{\pi}$, which we assume has full support ($\pi_i > 0$ for all i). Let us define the weighted inner product

$$\langle \mathbf{u}, \mathbf{v} \rangle_{\boldsymbol{\pi}} = \sum_i \frac{u_i v_i}{\pi_i}, \quad (47)$$

and let us order the eigenvalues of $\hat{\mathcal{L}}^{\text{cl}}$ as

$$0 = \Re(\lambda_1) > \Re(\lambda_2) \geq \Re(\lambda_3) \geq \dots. \quad (48)$$

We further assume the spectrum to not be degenerate for λ_j , with $j = 2$ in the case of athermality (see [Sec. 4.3](#)) and j corresponding to the slowest asymmetric eigenmode for asymmetry (see [Sec. 4.4](#)). Denoting the corresponding eigenmodes by $\mathbf{r}_k$ (orthonormal under $\langle \cdot, \cdot \rangle_{\boldsymbol{\pi}}$),

$$\hat{\mathcal{L}}^{\text{cl}} \mathbf{r}_k = \lambda_k \mathbf{r}_k, \quad \langle \mathbf{r}_k, \mathbf{r}_\ell \rangle_{\boldsymbol{\pi}} = \delta_{k\ell}. \quad (49)$$

In this eigenbasis an initial probability distribution $\mathbf{p}$ evolves as

$$\mathbf{p}(t) = \boldsymbol{\pi} + \sum_{k \geq 2} a_k \mathbf{r}_k e^{\lambda_k t}, \quad a_k = \langle \mathbf{p}, \mathbf{r}_k \rangle_{\boldsymbol{\pi}}, \quad (50)$$

so that $\delta \mathbf{p}(t) = \mathbf{p}(t) - \boldsymbol{\pi} = \sum_{k \geq 2} a_k \mathbf{r}_k e^{\lambda_k t}$. For small deviations one finds

$$D(\mathbf{p}(t) \parallel \boldsymbol{\pi}) = \sum_k p_k(t) \ln \frac{p_k(t)}{\pi_k} \sim \frac{1}{2} \sum_k \frac{(\delta p_k(t))^2}{\pi_k}. \quad (51)$$

As $t \rightarrow \infty$, as noted in Ref. [\[345\]](#), only the slowest decaying mode ($j = 2$) survives, $\delta p_k \sim a_2 r_{2,k} e^{\lambda_2 t}$, giving

$$D(\mathbf{p}(t) \parallel \boldsymbol{\pi}) \sim \frac{1}{2} a_2^2 e^{2\Re(\lambda_2)t} \sum_k \frac{r_{2,k}^2}{\pi_k}. \quad (4.19)$$

Replacing $\boldsymbol{\pi}$ with $\hat{\mathcal{G}}\mathbf{p}(t)$ allows us to extend this to the asymptotic behaviour of the relative entropy of asymmetry.

F.2 Proof of the Asymptotic Decay of the Quantum Relative Entropy

In both the thermalisation example of [Sec. 4.3.3](#) and the symmetry-breaking scenario of [Sec. 4.4](#), our main monotone

$$M(\hat{\rho}(t)) = S(\hat{\rho}(t) \parallel \hat{\pi}_\beta) \quad \text{or} \quad S(\hat{\rho}(t) \parallel \mathcal{G}[\hat{\rho}(t)]) \quad (52)$$

is the quantum relative entropy between the evolving state and its target (the relative entropy of athermality and of asymmetry, respectively).

At long times, its decay is governed by the slowest eigenmode of the Lindbladian $\mathcal{L}$ having nonzero overlap with the initial state, following [Eq. \(4.26\)](#) for the thermal Mpemba effect and [Eq. \(4.66\)](#) (in the case of the $\mathbb{Z}_4$ example studied in [Sec. 4.4.4](#)) for the asymmetry Mpemba effect.

Let now $\mathcal{L}$ be diagonalisable. Then there exist right eigenmodes $\{\hat{r}_k\}$ and left eigenmodes $\{\hat{\ell}_k\}$ satisfying

$$\mathcal{L}[\hat{r}_k] = \lambda_k \hat{r}_k, \quad \mathcal{L}^\dagger[\hat{\ell}_k] = \lambda_k^* \hat{\ell}_k, \quad (53)$$

with the biorthonormality condition

$$\text{Tr}[\hat{\ell}_k^\dagger \hat{r}_\ell] = \delta_{k\ell}, \quad (54)$$

and corresponding spectrum ordered so that

$$0 = \Re(\lambda_1) > \Re(\lambda_2) \geq \Re(\lambda_3) \geq \dots. \quad (55)$$

As in [Sec. F.1](#), we also assume the slowest eigenmode associated to eigenvalue λ_j to be non-degenerate (for athermality $j = 1$, for asymmetry the index j is chosen among eigenmodes in $\mu \neq 0$ -sectors as in the main text).

We now sketch the proof in three steps.

Spectral expansion. Analogously to [\(50\)](#), any initial state decomposes as

$$\hat{\rho}(t) = \hat{\pi} + \sum_{k \geq 2} c_k \hat{r}_k e^{\lambda_k t}, \quad c_k = \text{Tr}[\hat{\ell}_k^\dagger \hat{\rho}]. \quad (56)$$

Quadratic approximation of S . Writing $\hat{\rho} = \hat{\pi} + \delta\hat{\rho}$ with $\|\delta\hat{\rho}\| \ll 1$, one shows via the integral representation of the logarithm that

$$S(\hat{\rho} \parallel \hat{\pi}) = \text{Tr}[\hat{\rho}(\ln \hat{\rho} - \ln \hat{\pi})] \approx \frac{1}{2} \text{Tr}[\delta\hat{\rho} \hat{\pi}^{-1} \delta\hat{\rho}]. \quad (57)$$

Dominance of the slowest eigenmode. At late times $\delta\hat{\rho}(t) \approx c_j \hat{r}_j e^{\lambda_j t}$. If $\lambda_j \in \mathbb{C} \setminus \mathbb{R}$, its conjugate partner λ_j^* must also be included, and one finds

$$\delta\hat{\rho}(t) \approx c_j \hat{r}_j e^{\lambda_j t} + c_j^* \hat{r}_j^\dagger e^{\lambda_j^* t}. \quad (58)$$

Substituting into (57) gives

$$\begin{aligned} & \text{Tr}[\delta\hat{\rho} \hat{\pi}^{-1} \delta\hat{\rho}] \\ & \approx \begin{cases} |c_j|^2 e^{2\Re(\lambda_j)t} \text{Tr}[\hat{r}_j^\dagger \hat{\pi}^{-1} \hat{r}_j], & \lambda_j \in \mathbb{R}, \\ 2|c_j|^2 e^{2\Re(\lambda_j)t} \Re(\text{Tr}[\hat{r}_j^\dagger \hat{\pi}^{-1} \hat{r}_j]), & \lambda_j \notin \mathbb{R}. \end{cases} \end{aligned} \quad (59)$$

Hence,

$$S(\hat{\rho}(t) \parallel \hat{\pi}) \sim \begin{cases} \frac{1}{2} |c_j|^2 e^{2\Re(\lambda_j)t} \text{Tr}[\hat{r}_j^\dagger \hat{\pi}^{-1} \hat{r}_j], & \lambda_j \in \mathbb{R}, \\ |c_j|^2 e^{2\Re(\lambda_j)t} \Re(\text{Tr}[\hat{r}_j^\dagger \hat{\pi}^{-1} \hat{r}_j]), & \lambda_j \notin \mathbb{R}, \end{cases} \quad (60)$$

as discussed in Ref. [346]. Note that this can be straightforwardly extended to $S(\hat{\rho}(t) \parallel \mathcal{G}[\hat{\rho}(t)])$ by replacing $\hat{\pi}$ with $\mathcal{G}[\hat{\rho}(t)]$ in the proof.

G Self-dual kicked Ising gate

We relate the standard kicked-Ising (KI) definition of the self-dual kicked Ising (SDKI) model to the two-site brickwork gate $\hat{U}_{\text{SDKI}}(h)$ and record two exact two-qubit generators with $\exp[-i\hat{H}(h)] = \hat{U}_{\text{SDKI}}(h)$: the principal-branch six-string logarithm $\hat{H}_{\text{pr}}(h)$ and a four-string (non-principal) logarithm $\hat{H}_4(h)$. All identities are exact for the 4×4 gate.

G.1 Properties of the SDKI two-site gate

Bertini-Kos-Prosen define the kicked-Ising drive (period set to unity) as

$$\hat{H}_{\text{KI}}(t) = \hat{H}_I + \sum_{m \in \mathbb{Z}} \delta(t - m) \hat{H}_K, \quad (61)$$

with

$$\hat{H}_I = \frac{J}{2} \sum_{j=0}^{2L-1} \hat{\sigma}_j^z \hat{\sigma}_{j+1}^z + \frac{h}{2} \sum_{j=0}^{2L-1} \hat{\sigma}_j^z, \quad \hat{H}_K = \frac{b}{2} \sum_{j=0}^{2L-1} \hat{\sigma}_j^x, \quad (62)$$

and periodic boundary conditions. The Floquet operator is

$$\hat{U}_{\text{KI}} = \mathcal{T} \exp\left(-i \int_0^1 dt \hat{H}_{\text{KI}}(t)\right) = e^{-i\hat{H}_K} e^{-i\hat{H}_I}. \quad (63)$$

Splitting $\hat{H}_I$ into commuting even/odd bond layers yields a brickwork circuit; a convenient local two-site gate is

$$\hat{U}_{\text{KI}}^{(\text{local})} = e^{-iJ\hat{\sigma}^z\hat{\sigma}^z} \left(e^{-ih\hat{\sigma}^z} e^{-ib\hat{\sigma}^x} e^{-ih\hat{\sigma}^z} \otimes e^{-ib\hat{\sigma}^x} \right) e^{-iJ\hat{\sigma}^z\hat{\sigma}^z}. \quad (64)$$

At the *self-dual points*

$$|J| = |b| = \frac{\pi}{4}, \quad (65)$$

this brickwork circuit is dual-unitary and defines the SDKI model. A standard two-site representative is [330]

$$\hat{U}_{\text{SDKI}}(h) = \left(e^{-ih\hat{\sigma}^z} e^{-ib\hat{\sigma}^x} \otimes e^{-ib\hat{\sigma}^x} \right) e^{-iJ\hat{\sigma}^y\hat{\sigma}^y} e^{-iJ\hat{\sigma}^z\hat{\sigma}^z} \left(e^{-ih\hat{\sigma}^z} \otimes \hat{I} \right). \quad (66)$$

In the main text we adopt the gauge where the entangling block is

$$\hat{V}[0] = \exp\left[-i\frac{\pi}{4}(\hat{\sigma}^y\hat{\sigma}^y + \hat{\sigma}^z\hat{\sigma}^z)\right], \quad (67)$$

so that $\hat{U}_{\text{SDKI}}(h)$ is a dual-unitary gate for all h .

Expanding the single-qubit rotations (and using $[\hat{\sigma}^y\hat{\sigma}^y, \hat{\sigma}^z\hat{\sigma}^z] = 0$) gives the exact six-string Pauli decomposition

$$\hat{U}_{\text{SDKI}}(h) = \frac{1}{2} \left[-i(\hat{\sigma}^x\hat{I} - \hat{\sigma}^y\hat{\sigma}^y) - i\cos(2h)(\hat{I}\hat{\sigma}^x - \hat{\sigma}^z\hat{\sigma}^z) + \sin(2h)(\hat{\sigma}^z\hat{\sigma}^x - \hat{I}\hat{\sigma}^z) \right]. \quad (68)$$

Introduce three Hermitian involutions and abbreviations

$$\hat{g}_1 = \frac{\hat{\sigma}^x\hat{I} - \hat{\sigma}^y\hat{\sigma}^y}{\sqrt{2}}, \quad \hat{g}_2 = \frac{\hat{I}\hat{\sigma}^x - \hat{\sigma}^z\hat{\sigma}^z}{\sqrt{2}}, \quad \hat{g}_3 = \frac{\hat{\sigma}^z\hat{\sigma}^x - \hat{I}\hat{\sigma}^z}{\sqrt{2}}, \quad (69)$$

$$c := \cos(2h), \quad s := \sin(2h), \quad r := \sqrt{1+c^2}. \quad (70)$$

A direct Pauli-algebra computation yields

$$\hat{g}_1^2 = \hat{g}_2^2 = \hat{g}_3^2 = \hat{I}, \quad \{\hat{g}_1, \hat{g}_2\} = 0, \quad [\hat{g}_3, \hat{g}_1] = [\hat{g}_3, \hat{g}_2] = 0. \quad (71)$$

In this basis Eq. (68) reads

$$\hat{U}_{\text{SDKI}}(h) = \frac{1}{\sqrt{2}} \left(s \hat{g}_3 - i \hat{g}_1 - ic \hat{g}_2 \right) = \frac{1}{\sqrt{2}} \left(s \hat{g}_3 - ir \hat{g}_\perp \right), \quad \hat{g}_\perp = \frac{\hat{g}_1 + c \hat{g}_2}{r}. \quad (72)$$

By Eq. (71), $\hat{g}_\perp$ is Hermitian and satisfies $\hat{g}_\perp^2 = \hat{I}$, and moreover $[\hat{g}_\perp, \hat{g}_3] = 0$. Hence all logarithms reduce to the two commuting sectors of $\hat{g}_3$ with projectors

$$\hat{P}_\pm = \frac{\hat{I} \pm \hat{g}_3}{2}. \quad (73)$$

Restricting Eq. (72) to the $\hat{g}_3 = \pm 1$ sectors gives

$$\hat{U}_\pm = \hat{P}_\pm \hat{U}_{\text{SDKI}} \hat{P}_\pm = \frac{1}{\sqrt{2}} \left(\pm s \hat{I} - ir \hat{g}_\perp \right) = \cos \phi_\pm \hat{I} - i \sin \phi_\pm \hat{g}_\perp, \quad (74)$$

with $\cos \phi_\pm = \frac{\pm s}{\sqrt{2}}$, $\sin \phi_\pm = \frac{r}{\sqrt{2}}$. We choose

$$\theta(h) = \phi_+ = \text{atan2}(r, s) \in (0, \pi), \quad \phi_- = \text{atan2}(r, -s) = \pi - \theta(h), \quad (75)$$

so that $\phi_\pm \in (0, \pi)$ for all h (away from the measure-zero points where $\sin(2h) = 0$). This parametrisation will be the input for the logarithms below.

G.2 Connection to the four-term logarithm

Since $[\hat{g}_\perp, \hat{g}_3] = 0$ and $\hat{g}_\perp^2 = \hat{g}_3^2 = \hat{I}$, eigenphases may be shifted by 2π on individual joint eigenspaces of $(\hat{g}_3, \hat{g}_\perp)$ without changing $\exp[-i\hat{H}]$.

Define the joint eigenprojectors

$$\hat{\Pi}_{p,q} = \frac{1}{4} (\hat{I} + p \hat{g}_3) (\hat{I} + q \hat{g}_\perp), \quad p, q \in \{+1, -1\}. \quad (76)$$

Starting from the principal generator $\hat{H}_{\text{pr}}(h)$ in Sec. G.3, define the branch-shifted generator

$$\hat{H}_4(h) = \hat{H}_{\text{pr}}(h) + 2\pi \hat{\Pi}_{-, -}. \quad (77)$$

Using $\hat{\Pi}_{-, -} = \frac{1}{4} (\hat{I} - \hat{g}_3 - \hat{g}_\perp + \hat{g}_\perp \hat{g}_3)$ and the explicit g -basis form of $\hat{H}_{\text{pr}}(h)$ (cf. Eq. (86) below), the $\hat{g}_\perp$ term cancels *exactly* and one obtains the compact expression

$$\hat{H}_4(h) = \frac{\pi}{2} \hat{I} - \frac{\pi}{2} \hat{g}_3 + \theta(h) \hat{g}_\perp \hat{g}_3. \quad (78)$$

Since $\hat{g}_3$ is a difference of two one-/two-body strings and $\hat{g}_\perp \hat{g}_3$ reduces to two strings, this logarithm has strictly fewer Pauli strings than the principal one.

Expanding back to Pauli strings using Eq. (69) and the identities

$$\hat{g}_1 \hat{g}_3 = -\hat{\sigma}^x \hat{\sigma}^z, \quad \hat{g}_2 \hat{g}_3 = +\hat{\sigma}^z \hat{I}, \quad (79)$$

yields the advertised four-string form

$$\hat{H}_4(h) = \frac{\pi}{2} \hat{I} \hat{I} + \frac{\pi}{2\sqrt{2}} (\hat{I} \hat{\sigma}^z - \hat{\sigma}^z \hat{\sigma}^x) - \frac{\theta(h)}{r} \hat{\sigma}^x \hat{\sigma}^z + \frac{c\theta(h)}{r} \hat{\sigma}^z \hat{I}. \quad (80)$$

Because $[\hat{g}_\perp, \hat{g}_3] = 0$, one may work in the $\hat{g}_3 = \pm 1$ sectors. From Eq. (78) one has

$$\hat{H}_4|_{\hat{g}_3=+1} = \theta(h) \hat{g}_\perp, \quad \hat{H}_4|_{\hat{g}_3=-1} = \pi \hat{I} - \theta(h) \hat{g}_\perp. \quad (81)$$

Exponentiating and using $\hat{g}_\perp^2 = \hat{I}$ gives

$$e^{-i\hat{H}_4}|_{\hat{g}_3=+1} = \cos \theta \hat{I} - i \sin \theta \hat{g}_\perp, \quad (82)$$

$$e^{-i\hat{H}_4}|_{\hat{g}_3=-1} = e^{-i\pi} (\cos \theta \hat{I} + i \sin \theta \hat{g}_\perp) = -\cos \theta \hat{I} - i \sin \theta \hat{g}_\perp. \quad (83)$$

Using $\cos \theta = s/\sqrt{2}$ and $\sin \theta = r/\sqrt{2}$ from Eqs. (74) and (75), these are precisely the blocks Eq. (74). Recombining sectors therefore yields $e^{-i\hat{H}_4(h)} = \hat{U}_{\text{SDKI}}(h)$.

G.3 Connection to the principal six-term logarithm

We now record the principal-branch Floquet generator $\hat{H}_{\text{pr}}(h)$ defined by

$$\hat{U}_{\text{SDKI}}(h) = \exp[-i\hat{H}_{\text{pr}}(h)], \quad \text{spec}(\hat{H}_{\text{pr}}(h)) \subset (-\pi, \pi]. \quad (84)$$

From Eq. (74) each sector is an $\text{SU}(2)$ rotation with angle $\phi_\pm \in (0, \pi)$, hence

$$\hat{H}_{\text{pr}}(h) = \phi_+ \hat{g}_\perp \hat{P}_+ + \phi_- \hat{g}_\perp \hat{P}_-. \quad (85)$$

Using $\hat{P}_\pm = (\hat{I} \pm \hat{g}_3)/2$ and $\phi_- = \pi - \phi_+ = \pi - \theta(h)$ gives the closed g -basis form

$$\begin{aligned} \hat{H}_{\text{pr}}(h) &= \hat{g}_\perp \left(A \hat{I} + B \hat{g}_3 \right), \quad A = \frac{\phi_+ + \phi_-}{2} = \frac{\pi}{2}, \\ B &= \frac{\phi_+ - \phi_-}{2} = \theta(h) - \frac{\pi}{2} = \frac{2\theta(h) - \pi}{2}. \end{aligned} \quad (86)$$

Expanding $\hat{g}_\perp = (\hat{g}_1 + c\hat{g}_2)/r$ and using Eq. (69) and the products $\hat{g}_1 \hat{g}_3 = -\hat{\sigma}^x \hat{\sigma}^z$, $\hat{g}_2 \hat{g}_3 = +\hat{\sigma}^z \hat{I}$ yields a six-string Pauli expansion. Writing

$$\begin{aligned} \hat{H}_{\text{pr}}(h) &= a_{XI}(h) \hat{\sigma}^x \hat{I} + a_{YY}(h) \hat{\sigma}^y \hat{\sigma}^y + a_{IX}(h) \hat{I} \hat{\sigma}^x \\ &\quad + a_{ZZ}(h) \hat{\sigma}^z \hat{\sigma}^z + a_{XZ}(h) \hat{\sigma}^x \hat{\sigma}^z + a_{ZI}(h) \hat{\sigma}^z \hat{I}, \end{aligned} \quad (87)$$

with $c = \cos(2h)$, $r = \sqrt{1 + c^2}$, and $\theta(h) = \text{atan2}(r, \sin(2h))$, one finds

$$a_{XI}(h) = \frac{\pi}{2\sqrt{2}r}, \quad a_{YY}(h) = -\frac{\pi}{2\sqrt{2}r}, \quad (88)$$

$$a_{IX}(h) = \frac{\pi c}{2\sqrt{2}r}, \quad a_{ZZ}(h) = -\frac{\pi c}{2\sqrt{2}r}, \quad (89)$$

$$a_{XZ}(h) = -\frac{2\theta(h) - \pi}{2r}, \quad a_{ZI}(h) = \frac{c[2\theta(h) - \pi]}{2r}. \quad (90)$$

The difference between $\hat{H}_4(h)$ and $\hat{H}_{\text{pr}}(h)$ is a 2π branch shift on a *single* joint eigenspace. Indeed, using [Eq. \(86\)](#) and $\hat{\Pi}_{-,-} = \frac{1}{4}(\hat{I} - \hat{g}_3 - \hat{g}_\perp + \hat{g}_\perp \hat{g}_3)$, one computes

$$\hat{H}_{\text{pr}}(h) + 2\pi\hat{\Pi}_{-,-} = \hat{g}_\perp \left(\frac{\pi}{2}\hat{I} + \left(\theta - \frac{\pi}{2} \right) \hat{g}_3 \right) + \frac{\pi}{2} \left(\hat{I} - \hat{g}_3 - \hat{g}_\perp + \hat{g}_\perp \hat{g}_3 \right) \quad (91)$$

$$= \frac{\pi}{2}\hat{I} - \frac{\pi}{2}\hat{g}_3 + \theta \hat{g}_\perp \hat{g}_3 = \hat{H}_4(h), \quad (92)$$

which is exactly [Eq. \(78\)](#). In other words, according to [Eq. \(77\)](#), $\hat{H}_4(h)$ and $\hat{H}_{\text{pr}}(h)$ generate the same gate because they differ by 2π on a single eigenspace, and the reduction from six to four Pauli strings is the explicit cancellation of the $\hat{g}_\perp$ term in this branch choice.

List of Acronyms

- COE** circular orthogonal ensemble. 20, 46, 48
- CPTP** completely positive and trace-preserving. 90, 91, 99, 102, 104, 113, 122
- CUE** circular unitary ensemble. 20, 40
- DOS** density of states. 4, 8, 15, 16, 19, 20, 23, 24, 28–31, 48, 49, 53, 54, 60–63, 131, 170
- DQC1** deterministic quantum computation with one clean qubit. 49, 55–57, 178, 179
- DQT** dynamical quantum typicality. 52
- DU** dual unitary. 3, 34, 35, 37–39, 81, 83, 150, 152, 156, 157, 159–167, 184, 185
- ED** exact diagonalisation. 23, 28–31, 40, 60, 61, 63, 161
- ETH** eigenstate thermalisation hypothesis. 4, 12, 13, 15, 19–22, 40–42, 98, 131
- GGE** generalised Gibbs ensemble. 13, 42, 68, 82
- GHD** generalised hydrodynamics. 13
- GHZ** Greenberger-Horne-Zeilinger. 58
- GOE** Gaussian orthogonal ensemble. 20
- GUE** Gaussian unitary ensemble. 20
- KDE** kernel density estimate. 63
- KL** Kullback-Leibler. 91, 93, 95, 96, 108
- KMS** Kubo-Martin-Schwinger. 25, 26
- KPM** kernel polynomial method. 16, 23, 28, 31, 49, 51–53, 55, 60, 62, 170, 178
- KPZ** Kardar-Parisi-Zhang. 32, 33, 70, 72
- LCU** linear combination of unitaries. 176
- LDOS** local density of states. 4, 8, 15, 16, 19, 23, 48, 54, 131, 135
- LOCC** local operations and classical communication. 90

- LOE** local-operator entanglement. 131
- MBL** many-body localisation. 13, 14, 20, 22, 35, 38, 40, 136
- MEO** maximum eigenprojector overlap. 5, 136, 137, 151, 152, 156, 157, 160–167
- MEU** maximally ergodic unitary. 137, 138, 142, 143, 146, 148, 149
- MSD** mean-squared displacement. 9
- NISQ** noisy intermediate-scale quantum. 2, 178
- OBC** open boundary condition. 12, 45, 46, 111, 122
- OGH** operator growth hypothesis. 131, 136–138, 144, 149, 150, 166, 167, 171–173
- OTOC** out-of-time-order correlator. 4, 130, 131, 135
- PBC** periodic boundary condition. 12, 28, 30, 32, 62, 101, 122, 158
- POLFED** polynomially filtered exact diagonalisation method. 16, 38, 39, 49, 50, 83
- QCA** quantum cellular automaton. 164, 165, 167
- QSP** quantum signal processing. 176, 178
- QSVT** quantum singular value transformation. 57, 170, 176, 178, 179
- RMT** random matrix theory. 4, 131, 136
- RT** resource theory. 87, 89, 90, 92, 122
- SDKI** self-dual kicked Ising. 163, 171, 189
- SFF** spectral form factor. 15
- ST** Suzuki-Trotter. 32, 58–60, 62, 143, 153, 160, 164, 165
- SVD** singular value decomposition. 50, 51
- TFIM** transverse-field Ising model. 112, 120
- TO** thermal operation. 90, 92, 93, 103, 120
- ZDZ** Žnidarič-Duh-Zadnik. 180–182

References

- [1] A. Summer, C. Chiaracane, M. T. Mitchison, et al. “Calculating the Many-Body Density of States on a Digital Quantum Computer”. In: *Phys. Rev. Research* 6.1 (Jan. 2024), p. 013106.
- [2] A. Summer, A. Nico-Katz, S. Dooley, et al. “Anomalous transport in U(1)-symmetric quantum circuits”. In: *npj Quantum Information* 12.1, 32 (Jan. 2026).
- [3] A. Summer, M. Moroder, L. P. Bettmann, et al. “Resource-theoretical unification of Mpemba effects: classical and quantum”. In: *Physical Review X* 16.1 (Mar. 2026), p. 011065.
- [4] S. Aditya, A. Summer, P. Sierant, et al. *Mpemba Effects in Quantum Complexity*. 2025.
- [5] A. Summer, S. Dooley, and J. Goold. “Prethermalisation induced by an impurity in a U(1)-symmetric quantum circuit”. Manuscript in preparation. 2026.
- [6] A. Summer, H. K. Yadalam, and M. Mitchison. “Krylov complexity of Floquet generators of quantum circuits”. Manuscript in preparation. 2026.
- [7] K. Blekos, D. Brand, A. Ceschini, et al. “A review on Quantum Approximate Optimization Algorithm and its variants”. In: *Physics Reports* 1068 (June 2024), pp. 1–66.
- [8] J. P. Dowling and G. J. Milburn. *Quantum Technology: The Second Quantum Revolution*. 2002.
- [9] C. Monroe and J. Kim. “Scaling the Ion Trap Quantum Processor”. In: *Science* 339.6124 (Mar. 2013), pp. 1164–1169.
- [10] M. Saffman, T. G. Walker, and K. Mølmer. “Quantum information with Rydberg atoms”. In: *Rev. Mod. Phys.* 82 (3 Aug. 2010), pp. 2313–2363.
- [11] I. Bloch, J. Dalibard, and W. Zwerger. “Many-body physics with ultracold gases”. In: *Rev. Mod. Phys.* 80 (3 July 2008), pp. 885–964.
- [12] M. H. Devoret and R. J. Schoelkopf. “Superconducting Circuits for Quantum Information: An Outlook”. In: *Science* 339.6124 (Mar. 2013), pp. 1169–1174.
- [13] J. Preskill. “Quantum Computing in the NISQ era and beyond”. In: *Quantum* 2 (Aug. 2018), p. 79.
- [14] C. L. Degen, F. Reinhard, and P. Cappellaro. “Quantum sensing”. In: *Reviews of Modern Physics* 89.3 (July 2017).
- [15] N. Gisin, G. Ribordy, W. Tittel, et al. “Quantum cryptography”. In: *Reviews of Modern Physics* 74.1 (Mar. 2002), pp. 145–195.
- [16] P. Shor. “Algorithms for quantum computation: discrete logarithms and factoring”. In: *Proceedings 35th Annual Symposium on Foundations of Computer Science*. SFCS-94. IEEE Comput. Soc. Press, pp. 124–134.
- [17] R. P. Feynman. “Simulating physics with computers”. In: *International Journal of Theoretical Physics* 21.6-7 (June 1982), pp. 467–488.
- [18] S. Lloyd. “Universal Quantum Simulators”. In: *Science* 273.5278 (1996), pp. 1073–1078.

-
- [19] I. M. Georgescu, S. Ashhab, and F. Nori. “Quantum simulation”. In: *Reviews of Modern Physics* 86.1 (Mar. 2014), pp. 153–185.
 - [20] I. H. Deutsch. “Harnessing the Power of the Second Quantum Revolution”. In: *PRX Quantum* 1 (2 Nov. 2020), p. 020101.
 - [21] F. Arute et al. “Quantum Supremacy Using a Programmable Superconducting Processor”. In: *Nature* 574 (2019), pp. 505–510.
 - [22] S. Coleman. *Quote: The career of a young theoretical physicist consists of treating the harmonic oscillator in ever-increasing levels of abstraction*. Anecdotal, cited in various sources. 1985.
 - [23] M. Bukov, L. D’Alessio, and A. Polkovnikov. “Universal high-frequency behavior of periodically driven systems: from dynamical stabilization to Floquet engineering”. In: *Advances in Physics* 64.2 (Mar. 2015), pp. 139–226.
 - [24] A. Eckardt. “Colloquium: Atomic Quantum Gases in Periodically Driven Optical Lattices”. In: *Reviews of Modern Physics* 89.1 (2017), p. 011004.
 - [25] S. Wolfram. “Statistical Mechanics of Cellular Automata”. In: *Reviews of Modern Physics* 55.3 (1983), pp. 601–644.
 - [26] S. Wolfram. *A New Kind of Science*. Wolfram Media, 2002.
 - [27] T. Nagatani. “The Physics of Traffic Jams”. In: *Reports on Progress in Physics* 65.9 (2002), pp. 1331–1386.
 - [28] B. Chopard and M. Droz. “Cellular Automata Modeling of Physical Systems”. In: *Cambridge University Press* (1998).
 - [29] M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information*. 10th Anniversary Edition. Cambridge University Press, 2010.
 - [30] J. Watrous. “On one-dimensional quantum cellular automata”. In: *Proceedings of IEEE 36th Annual Foundations of Computer Science*. SFCS-95. IEEE Comput. Soc. Press, pp. 528–537.
 - [31] B. Schumacher and R. F. Werner. *Reversible quantum cellular automata*. 2004.
 - [32] P. Arrighi. “An Overview of Quantum Cellular Automata”. In: *Natural Computing* 18 (2019), pp. 885–899.
 - [33] J. Watrous. *The theory of quantum information*. en. Cambridge, England: Cambridge University Press, Apr. 2018.
 - [34] T. Mori, T. Kuwahara, and K. Saito. “Rigorous Bound on Energy Absorption and Generic Relaxation in Periodically Driven Quantum Systems”. In: *Physical Review Letters* 116.12 (Mar. 2016).
 - [35] D. A. Abanin, W. De Roeck, W. W. Ho, et al. “Effective Hamiltonians, prethermalization, and slow energy absorption in periodically driven many-body systems”. In: *Physical Review B* 95.1 (Jan. 2017).
 - [36] D. V. Else, C. Monroe, C. Nayak, et al. “Discrete Time Crystals”. In: *Annual Review of Condensed Matter Physics* 11 (2020), pp. 467–499.
 - [37] M. S. Rudner, N. H. Lindner, E. Berg, et al. “Anomalous Edge States and the Bulk-Edge Correspondence for Periodically Driven Two-Dimensional Systems”. In: *Physical Review X* 3.3 (July 2013).

-
- [38] A. Nahum, J. Ruhman, S. Vijay, et al. “Quantum Entanglement Growth under Random Unitary Dynamics”. In: *Physical Review X* 7.3 (July 2017).
 - [39] A. Nahum, S. Vijay, and J. Haah. “Operator Spreading in Random Unitary Circuits”. In: *Physical Review X* 8.2 (Apr. 2018).
 - [40] B. Skinner, J. Ruhman, and A. Nahum. “Measurement-Induced Phase Transitions in the Dynamics of Entanglement”. In: *Phys. Rev. X* 9 (3 July 2019), p. 031009.
 - [41] B. Bertini, P. Kos, and T. Prosen. “Exact Correlation Functions for Dual-Unitary Lattice Models in $1 + 1$ Dimensions”. In: *Phys. Rev. Lett.* 123 (21 Nov. 2019), p. 210601.
 - [42] D. Forster. *Hydrodynamic Fluctuations, Broken Symmetry, and Correlation Functions*. CRC Press, Mar. 2018.
 - [43] B. Bertini, F. Heidrich-Meisner, C. Karrasch, et al. “Finite-temperature transport in one-dimensional quantum lattice models”. In: *Rev. Mod. Phys.* 93 (2 May 2021), p. 025003.
 - [44] G. Casati, F. Valz-Gris, and I. Guarneri. “On the connection between quantization of nonintegrable systems and statistical theory of spectra”. In: *Lettere al Nuovo Cimento* 28.8 (June 1980), pp. 279–282.
 - [45] F. Haake. *Quantum Signatures of Chaos*. Springer Berlin Heidelberg, 2010.
 - [46] L. D’Alessio, Y. Kafri, A. Polkovnikov, et al. “From quantum chaos and eigenstate thermalization to statistical mechanics and thermodynamics”. In: *Advances in Physics* 65.3 (May 2016), pp. 239–362.
 - [47] C. J. Turner, A. A. Michailidis, D. A. Abanin, et al. “Weak ergodicity breaking from quantum many-body scars”. In: *Nature Physics* 14.7 (May 2018), pp. 745–749.
 - [48] J. Maldacena, S. H. Shenker, and D. Stanford. “A bound on chaos”. In: *Journal of High Energy Physics* 2016.8 (Aug. 2016).
 - [49] C. von Keyserlingk, T. Rakovszky, F. Pollmann, et al. “Operator Hydrodynamics, OTOCs, and Entanglement Growth in Systems without Conservation Laws”. In: *Physical Review X* 8.2 (Apr. 2018).
 - [50] R. Kubo. “Statistical-Mechanical Theory of Irreversible Processes. I. General Theory and Simple Applications to Magnetic and Conduction Problems”. In: *Journal of the Physical Society of Japan* 12.6 (June 1957), pp. 570–586.
 - [51] A. Einstein. “Über die von der molekularkinetischen Theorie der Wärme geforderte Bewegung von in ruhenden Flüssigkeiten suspendierten Teilchen”. In: *Annalen der Physik* 322.8 (Jan. 1905), pp. 549–560.
 - [52] J. B. J. Fourier. *The Analytical Theory of Heat*. Ed. by A. Freeman. Cambridge University Press, July 1822.
 - [53] L. P. Kadanoff and P. C. Martin. “Hydrodynamic equations and correlation functions”. In: *Annals of Physics* 24 (Oct. 1963), pp. 419–469.
 - [54] G. D. Mahan. *Many-Particle Physics*. Springer US, 2000.
 - [55] P. Mazur. “Non-ergodicity of phase functions in certain systems”. In: *Physica* 43.4 (Sept. 1969), pp. 533–545.
 - [56] M. Suzuki. “Ergodicity, constants of motion, and bounds for susceptibilities”. In: *Physica* 51.2 (Jan. 1971), pp. 277–291.

-
- [57] C. Gogolin and J. Eisert. “Equilibration, thermalisation, and the emergence of statistical mechanics in closed quantum systems”. In: *Reports on Progress in Physics* 79.5 (Apr. 2016), p. 056001.
 - [58] R. Nandkishore and D. A. Huse. “Many-Body Localization and Thermalization in Quantum Statistical Mechanics”. In: *Annual Review of Condensed Matter Physics* 6.1 (Mar. 2015), pp. 15–38.
 - [59] M. Rigol and B. S. Shastry. “Drude weight in systems with open boundary conditions”. In: *Phys. Rev. B* 77 (16 Apr. 2008), p. 161101.
 - [60] R. Steinigeweg, J. Gemmer, and W. Brenig. “Spin-Current Autocorrelations from Single Pure-State Propagation”. In: *Physical Review Letters* 112.12 (Mar. 2014).
 - [61] S. Pilatowsky-Cameo and S. Choi. “Quantum thermalization must occur in translation-invariant systems at high temperature”. In: *Nature Communications* (Dec. 2025).
 - [62] S. Pilatowsky-Cameo, I. Marvian, S. Choi, et al. “Hilbert-Space Ergodicity in Driven Quantum Systems: Obstructions and Designs”. In: *Physical Review X* 14.4 (Dec. 2024).
 - [63] A. L. Shaw, D. K. Mark, J. Choi, et al. “Experimental Signatures of Hilbert-Space Ergodicity: Universal Bitstring Distributions and Applications in Noise Learning”. In: *Physical Review X* 15.3 (July 2025).
 - [64] O. A. Castro-Alvaredo, B. Doyon, and T. Yoshimura. “Emergent Hydrodynamics in Integrable Quantum Systems Out of Equilibrium”. In: *Physical Review X* 6.4 (Dec. 2016).
 - [65] B. Bertini, M. Collura, J. De Nardis, et al. “Transport in Out-of-Equilibrium XXZ Chains: Exact Profiles of Charges and Currents”. In: *Phys. Rev. Lett.* 117 (20 Nov. 2016), p. 207201.
 - [66] M. Schemmer, I. Bouchoule, B. Doyon, et al. “Generalized Hydrodynamics on an Atom Chip”. In: *Physical Review Letters* 122.9 (Mar. 2019).
 - [67] N. Malvania, Y. Zhang, Y. Le, et al. “Generalized hydrodynamics in strongly interacting 1D Bose gases”. In: *Science* 373.6559 (Sept. 2021), pp. 1129–1133.
 - [68] I. Bouchoule and J. Dubail. “Generalized hydrodynamics in the one-dimensional Bose gas: theory and experiments”. In: *Journal of Statistical Mechanics: Theory and Experiment* 2022.1 (Jan. 2022), p. 014003.
 - [69] A. Polkovnikov, K. Sengupta, A. Silva, et al. “Colloquium: Nonequilibrium dynamics of closed interacting quantum systems”. In: *Reviews of Modern Physics* 83.3 (Aug. 2011), pp. 863–883.
 - [70] J. Gemmer, R. Steinigeweg, and M. Michel. “Limited validity of the Kubo formula for thermal conduction in modular quantum systems”. In: *Phys. Rev. B* 73 (10 Mar. 2006), p. 104302.
 - [71] V. Oganesyan and D. A. Huse. “Localization of interacting fermions at high temperature”. In: *Phys. Rev. B* 75 (15 Apr. 2007), p. 155111.
 - [72] Y. Y. Atas, E. Bogomolny, O. Giraud, et al. “Distribution of the Ratio of Consecutive Level Spacings in Random Matrix Ensembles”. In: *Phys. Rev. Lett.* 110 (8 Feb. 2013), p. 084101.

-
- [73] M. Srednicki. “Chaos and quantum thermalization”. In: *Physical Review E* 50.2 (1994), pp. 888–901.
 - [74] T. Prosen and E. Ilievski. “Families of Quasilocal Conservation Laws and Quantum Spin Transport”. In: *Physical Review Letters* 111.5 (Aug. 2013).
 - [75] P. Reimann. “Foundation of Statistical Mechanics under Experimentally Realistic Conditions”. In: *Physical Review Letters* 101.19 (Nov. 2008).
 - [76] A. J. Short. “Equilibration of quantum systems and subsystems”. In: *New Journal of Physics* 13.5 (May 2011), p. 053009.
 - [77] N. Linden, S. Popescu, A. J. Short, et al. “Quantum mechanical evolution towards thermal equilibrium”. In: *Physical Review E* 79.6 (June 2009).
 - [78] F. J. Dyson. “The Threefold Way. Algebraic Structure of Symmetry Groups and Ensembles in Quantum Mechanics”. In: *Journal of Mathematical Physics* 3.6 (Nov. 1962), pp. 1199–1215.
 - [79] P. Ponte, Z. Papić, F. Huveneers, et al. “Many-Body Localization in Periodically Driven Systems”. In: *Phys. Rev. Lett.* 114 (14 Apr. 2015), p. 140401.
 - [80] J. M. Deutsch. “Quantum statistical mechanics in a closed system”. In: *Physical Review A* 43.4 (1991), pp. 2046–2049.
 - [81] J. M. Deutsch. “Eigenstate thermalization hypothesis”. In: *Reports on Progress in Physics* 81.8 (July 2018), p. 082001.
 - [82] H. Kim, T. N. Ikeda, and D. A. Huse. “Testing whether all eigenstates obey the eigenstate thermalization hypothesis”. In: *Phys. Rev. E* 90 (5 Nov. 2014), p. 052105.
 - [83] T. Kuwahara and K. Saito. “Eigenstate thermalization from the clustering property of correlation”. In: *Physical review letters* 124.20 (2020), p. 200604.
 - [84] D. N. Page. “Average entropy of a subsystem”. In: *Phys. Rev. Lett.* 71 (9 Aug. 1993), pp. 1291–1294.
 - [85] E. Bianchi and P. Dona. “Typical entanglement entropy in the presence of a center: Page curve and its variance”. In: *Physical Review D* 100.10 (2019), p. 105010.
 - [86] V. Balachandran, L. F. Santos, M. Rigol, et al. “Slow relaxation of out-of-time-ordered correlators in interacting integrable and nonintegrable spin- $\frac{1}{2}$ XYZ chains”. In: *Phys. Rev. B* 107 (23 June 2023), p. 235421.
 - [87] M. Hartmann, G. Mahler, and O. Hess. “Spectral Densities and Partition Functions of Modular Quantum systems as Derived from a Central Limit Theorem”. In: *Journal of Statistical Physics* 119.5-6 (June 2005), pp. 1139–1151.
 - [88] M. Will, T. A. Cochran, E. Rosenberg, et al. “Probing non-equilibrium topological order on a quantum processor”. In: *Nature* 645.8080 (Sept. 2025), pp. 348–353.
 - [89] X. Zhang, W. Jiang, J. Deng, et al. “Digital quantum simulation of Floquet symmetry-protected topological phases”. In: *Nature* 607.7919 (July 2022), pp. 468–473.
 - [90] T. N. Ikeda, S. Sugiura, and A. Polkovnikov. “Robust Effective Ground State in a Nonintegrable Floquet Quantum Circuit”. In: *Physical Review Letters* 133.3 (July 2024).

-
- [91] S. Bhattacharjee, S. Bandyopadhyay, and A. Polkovnikov. “Sharp detection of the onset of Floquet heating using eigenstate sensitivity”. In: *The European Physical Journal B* **97.10** (Oct. 2024).
 - [92] N. Keenan, N. F. Robertson, T. Murphy, et al. “Evidence of Kardar-Parisi-Zhang scaling on a digital quantum simulator”. In: *npj Quantum Information* **9.1** (July 2023).
 - [93] M. Žnidarič, U. Duh, and L. Zadnik. “Integrability is generic in homogeneous $U(1)$ -invariant nearest-neighbor qubit circuits”. In: *Physical Review B* **112.2** (July 2025).
 - [94] H. Bethe. “Zur Theorie der Metalle: I. Eigenwerte und Eigenfunktionen der linearen Atomkette”. In: *Zeitschrift für Physik* **71.3-4** (Mar. 1931), pp. 205–226.
 - [95] T. Prosen. “Open XXZ Spin Chain: Nonequilibrium Steady State and a Strict Bound on Ballistic Transport”. In: *Phys. Rev. Lett.* **106** (21 May 2011), p. 217206.
 - [96] M. Žnidarič, A. Scardicchio, and V. K. Varma. “Diffusive and Subdiffusive Spin Transport in the Ergodic Phase of a Many-Body Localizable System”. In: *Phys. Rev. Lett.* **117** (4 July 2016), p. 040601.
 - [97] A. Pal and D. A. Huse. “Many-body localization phase transition”. In: *Phys. Rev. B* **82** (17 Nov. 2010), p. 174411.
 - [98] D. J. Luitz, N. Laflorencie, and F. Alet. “Many-body localization edge in the random-field Heisenberg chain”. In: *Phys. Rev. B* **91** (8 Feb. 2015), p. 081103.
 - [99] D. A. Abanin, E. Altman, I. Bloch, et al. “Colloquium: Many-body localization, thermalization, and entanglement”. In: *Rev. Mod. Phys.* **91** (2 May 2019), p. 021001.
 - [100] A. Nico-Katz, A. Bayat, and S. Bose. “Information-theoretic memory scaling in the many-body localization transition”. In: *Phys. Rev. B* **105** (20 May 2022), p. 205133.
 - [101] A. Nico-Katz, A. Bayat, and S. Bose. “Memory hierarchy for many-body localization: Emulating the thermodynamic limit”. In: *Phys. Rev. Res.* **4** (3 July 2022), p. 033070.
 - [102] M. Serbyn, Z. Papić, and D. A. Abanin. “Quantum quenches in the many-body localized phase”. In: *Phys. Rev. B* **90** (17 Nov. 2014), p. 174302.
 - [103] M. Serbyn, Z. Papić, and D. A. Abanin. “Criterion for Many-Body Localization-Delocalization Phase Transition”. In: *Phys. Rev. X* **5** (4 Dec. 2015), p. 041047.
 - [104] V. Khemani, D. N. Sheng, and D. A. Huse. “Two Universality Classes for the Many-Body Localization Transition”. In: *Phys. Rev. Lett.* **119** (7 Aug. 2017), p. 075702.
 - [105] D. J. Luitz, I. M. Khaymovich, and Y. B. Lev. “Multifractality and its role in anomalous transport in the disordered XXZ spin-chain”. In: *SciPost Phys. Core* **2** (2020), p. 006.
 - [106] A. Nico-Katz and S. Bose. “Entanglement-complexity geometric measure”. In: *Phys. Rev. Res.* **5** (1 Jan. 2023), p. 013041.
 - [107] J. Colbois, F. Alet, and N. Laflorencie. “Interaction-Driven Instabilities in the Random-Field XXZ Chain”. In: *Phys. Rev. Lett.* **133** (11 Sept. 2024), p. 116502.
 - [108] M. Vanicat, L. Zadnik, and T. Prosen. “Integrable Trotterization: Local Conservation Laws and Boundary Driving”. In: *Phys. Rev. Lett.* **121** (3 July 2018), p. 030606.
 - [109] M. Ljubotina, L. Zadnik, and T. Prosen. “Ballistic Spin Transport in a Periodically Driven Integrable Quantum System”. In: *Phys. Rev. Lett.* **122** (15 Apr. 2019), p. 150605.
 - [110] M. Ljubotina, M. Žnidarič, and T. Prosen. “Kardar-Parisi-Zhang Physics in the Quantum Heisenberg Magnet”. In: *Phys. Rev. Lett.* **122** (21 May 2019), p. 210602.

-
- [111] R. Peierls. “Zur Theorie des Diamagnetismus von Leitungselektronen”. In: *Zeitschrift für Physik* 80.11-12 (1933), pp. 763–791.
 - [112] J. Zhang, J. Vala, S. Sastry, et al. “Geometric theory of nonlocal two-qubit operations”. In: *Physical Review A* 67.4 (Apr. 2003).
 - [113] Y. Bar Lev, G. Cohen, and D. R. Reichman. “Absence of Diffusion in an Interacting System of Spinless Fermions on a One-Dimensional Disordered Lattice”. In: *Phys. Rev. Lett.* 114 (10 Mar. 2015), p. 100601.
 - [114] C. Sünderhauf, D. Pérez-García, D. A. Huse, et al. “Localization with random time-periodic quantum circuits”. In: *Phys. Rev. B* 98 (13 Oct. 2018), p. 134204.
 - [115] A. Morningstar, L. Colmenarez, V. Khemani, et al. “Avalanches and many-body resonances in many-body localized systems”. In: *Phys. Rev. B* 105 (17 May 2022), p. 174205.
 - [116] D. Hahn and L. Colmenarez. “Absence of localization in weakly interacting Floquet circuits”. In: *Phys. Rev. B* 109 (9 Mar. 2024), p. 094207.
 - [117] O. Shtanko, D. S. Wang, H. Zhang, et al. “Uncovering local integrability in quantum many-body dynamics”. In: *Nature Communications* 16.1 (Mar. 2025).
 - [118] D. M. Long, P. J. D. Crowley, V. Khemani, et al. “Phenomenology of the Prethermal Many-Body Localized Regime”. In: *Phys. Rev. Lett.* 131 (10 Sept. 2023), p. 106301.
 - [119] V. Khemani, A. Vishwanath, and D. A. Huse. “Operator Spreading and the Emergence of Dissipative Hydrodynamics under Unitary Evolution with Conservation Laws”. In: *Physical Review X* 8.3 (Sept. 2018), p. 031057.
 - [120] T. Rakovszky, F. Pollmann, and C. W. von Keyserlingk. “Diffusive Hydrodynamics of Out-of-Time-Ordered Correlators with Charge Conservation”. In: *Phys. Rev. X* 8 (3 Sept. 2018), p. 031058.
 - [121] P. Sierant, M. Lewenstein, A. Scardicchio, et al. “Stability of many-body localization in Floquet systems”. In: *Phys. Rev. B* 107 (11 Mar. 2023), p. 115132.
 - [122] C. Jonay, J. F. Rodriguez-Nieva, and V. Khemani. “Slow thermalization and subdiffusion in $U(1)$ conserving Floquet random circuits”. In: *Phys. Rev. B* 109 (2 Jan. 2024), p. 024311.
 - [123] P. Sierant, M. Lewenstein, A. Scardicchio, et al. “Many-body localization in the age of classical computing*”. In: *Reports on Progress in Physics* 88.2 (Jan. 2025), p. 026502.
 - [124] R. K. Panda, A. Scardicchio, M. Schulz, et al. “Can we study the many-body localisation transition?” In: *Europhysics Letters* 128.6 (Feb. 2020), p. 67003.
 - [125] F. M. Dekking, C. Kraaikamp, H. P. Lopuhaä, et al. *A Modern Introduction to Probability and Statistics*. Springer London, 2005.
 - [126] P. Sierant, M. Lewenstein, and J. Zakrzewski. “Polynomially Filtered Exact Diagonalization Approach to Many-Body Localization”. In: *Phys. Rev. Lett.* 125 (15 Oct. 2020), p. 156601.
 - [127] W. De Roeck and F. Huveneers. “Stability and instability towards delocalization in many-body localization systems”. In: *Phys. Rev. B* 95 (15 Apr. 2017), p. 155129.
 - [128] T. Thiery, F. Huveneers, M. Müller, et al. “Many-Body Delocalization as a Quantum Avalanche”. In: *Phys. Rev. Lett.* 121 (14 Oct. 2018), p. 140601.

-
- [129] F. Alet and N. Laflorencie. “Many-body localization: An introduction and selected topics”. In: *Comptes Rendus. Physique* 19.6 (Apr. 2018), pp. 498–525.
 - [130] E. V. Doggen, I. V. Gornyi, A. D. Mirlin, et al. “Many-body localization in large systems: Matrix-product-state approach”. In: *Annals of Physics* 435 (Dec. 2021), p. 168437.
 - [131] M. Brenes, E. Mascarenhas, M. Rigol, et al. “High-temperature coherent transport in the XXZ chain in the presence of an impurity”. In: *Physical Review B* 98.23 (Dec. 2018).
 - [132] M. Brenes, T. LeBlond, J. Goold, et al. “Eigenstate Thermalization in a Locally Perturbed Integrable System”. In: *Phys. Rev. Lett.* 125 (7 Aug. 2020), p. 070605.
 - [133] D. J. Luitz. “Polynomial filter diagonalization of large Floquet unitary operators”. In: *SciPost Phys.* 11 (2021), p. 021.
 - [134] S. Popescu, A. J. Short, and A. Winter. “Entanglement and the foundations of statistical mechanics”. In: *Nature Physics* 2.11 (Oct. 2006), pp. 754–758.
 - [135] S. Goldstein, J. L. Lebowitz, R. Tumulka, et al. “Canonical Typicality”. In: *Physical Review Letters* 96.5 (Feb. 2006).
 - [136] A. Weiße, G. Wellein, A. Alvermann, et al. “The kernel polynomial method”. In: *Rev. Mod. Phys.* 78 (1 Mar. 2006), pp. 275–306.
 - [137] A. Hams and H. De Raedt. “Fast algorithm for finding the eigenvalue distribution of very large matrices”. In: *Physical Review E* 62.3 (Sept. 2000), pp. 4365–4377.
 - [138] M. Hutchinson. “A stochastic estimator of the trace of the influence matrix for laplacian smoothing splines”. In: *Communications in Statistics - Simulation and Computation* 19.2 (Jan. 1990), pp. 433–450.
 - [139] H. Avron and S. Toledo. “Randomized algorithms for estimating the trace of an implicit symmetric positive semi-definite matrix”. In: *Journal of the ACM* 58.2 (Apr. 2011), pp. 1–34.
 - [140] C. Bartsch and J. Gemmer. “Dynamical Typicality of Quantum Expectation Values”. In: *Physical Review Letters* 102.11 (Mar. 2009).
 - [141] R. Steinigeweg, J. Gemmer, and W. Brenig. “Spin and energy currents in integrable and nonintegrable spin-1/2 chains: A typicality approach to real-time autocorrelations”. In: *Physical Review B* 91.10 (Mar. 2015).
 - [142] S. Sugiura and A. Shimizu. “Thermal Pure Quantum States at Finite Temperature”. In: *Physical Review Letters* 108.24 (June 2012).
 - [143] J. E. Sobczyk and A. Roggero. “Spectral density reconstruction with Chebyshev polynomials”. In: *Phys. Rev. E* 105 (5 May 2022), p. 055310.
 - [144] D. Elliott and P. J. Taylor. “Polynomial approximation errors for functions of low-order continuity”. In: *Constructive Approximation* 7.1 (Dec. 1991), pp. 381–387.
 - [145] J. Emerson, Y. S. Weinstein, M. Saraceno, et al. “Pseudo-Random Unitary Operators for Quantum Information Processing”. In: *Science* 302.5653 (Dec. 2003), pp. 2098–2100.
 - [146] S. Boixo, S. V. Isakov, V. N. Smelyanskiy, et al. “Characterizing quantum supremacy in near-term devices”. In: *Nature Physics* 14.6 (Apr. 2018), pp. 595–600.

-
- [147] J. Richter and A. Pal. “Simulating Hydrodynamics on Noisy Intermediate-Scale Quantum Devices with Random Circuits”. In: *Phys. Rev. Lett.* **126** (23 June 2021), p. 230501.
 - [148] G. H. Low and I. L. Chuang. “Hamiltonian Simulation by Qubitization”. In: *Quantum* **3** (July 2019), p. 163.
 - [149] A. Gilyén, Y. Su, G. H. Low, et al. “Quantum singular value transformation and beyond: exponential improvements for quantum matrix arithmetics”. In: *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*. STOC ’19. ACM, June 2019, pp. 193–204.
 - [150] J. M. Martyn, Z. M. Rossi, A. K. Tan, et al. “Grand Unification of Quantum Algorithms”. In: *PRX Quantum* **2** (4 Dec. 2021), p. 040203.
 - [151] E. Knill and R. Laflamme. “Power of One Bit of Quantum Information”. In: *Physical Review Letters* **81.25** (1998), pp. 5672–5675.
 - [152] A. M. Childs, D. Maslov, Y. Nam, et al. “Toward the First Quantum Simulation with Quantum Speedup”. In: *Proceedings of the National Academy of Sciences* **115.38** (Sept. 2018), pp. 9456–9461.
 - [153] A. M. Childs, Y. Su, M. C. Tran, et al. “Theory of Trotter Error with Commutator Scaling”. In: *Phys. Rev. X* **11** (1 Feb. 2021), p. 011020.
 - [154] H. Wang, J. Nan, X. Qiu, et al. “Quantum Kernel Function Expansion for Thermal Quantum Ensemble”. In: *arXiv preprint arXiv:2202.01170* (2022).
 - [155] G. García-Pérez, M. A. Rossi, B. Sokolov, et al. “Learning to Measure: Adaptive Informationally Complete Generalized Measurements for Quantum Algorithms”. In: *PRX Quantum* **2** (4 Nov. 2021), p. 040342.
 - [156] H.-Y. Huang, R. Kueng, and J. Preskill. “Predicting many properties of a quantum system from very few measurements”. In: *Nature Physics* **16.10** (June 2020), pp. 1050–1057.
 - [157] Quantinuum. *H1-1 System*. 2022.
 - [158] D. Abanin, W. De Roeck, W. W. Ho, et al. “A Rigorous Theory of Many-Body Prethermalization for Periodically Driven and Closed Quantum Systems”. In: *Communications in Mathematical Physics* **354.3** (June 2017), pp. 809–827.
 - [159] T. Kuwahara, T. Mori, and K. Saito. “Floquet-Magnus theory and generic transient dynamics in periodically driven many-body quantum systems”. In: *Annals of Physics* **367** (Apr. 2016), pp. 96–124.
 - [160] W. W. Ho, T. Mori, D. A. Abanin, et al. “Quantum and classical Floquet prethermalization”. In: *Annals of Physics* **454** (July 2023), p. 169297.
 - [161] B. Bertini and M. Fagotti. “Pre-relaxation in weakly interacting models”. In: *Journal of Statistical Mechanics: Theory and Experiment* **2015.7** (July 2015), P07012.
 - [162] B. Bertini, F. H. L. Essler, S. Groha, et al. “Prethermalization and Thermalization in Models with Weak Integrability Breaking”. In: *Physical Review Letters* **115.18** (Oct. 2015).
 - [163] P. Reimann and L. Dabelow. “Typicality of Prethermalization”. In: *Physical Review Letters* **122.8** (Mar. 2019).

-
- [164] R. Steinigeweg and J. Gemmer. “Density dynamics in translationally invariant spin- $\frac{1}{2}$ chains at high temperatures: A current-autocorrelation approach to finite time and length scales”. In: *Phys. Rev. B* 80 (18 Nov. 2009), p. 184402.
 - [165] K. Mallayya, M. Rigol, and W. De Roeck. “Prethermalization and Thermalization in Isolated Quantum Systems”. In: *Phys. Rev. X* 9 (2 May 2019), p. 021027.
 - [166] U. Duh and M. Žnidarič. *Ruelle-Pollicott resonances of diffusive $U(1)$ -invariant qubit circuits*. 2025.
 - [167] T. Zhou and A. W. Harrow. “Maximal entanglement velocity implies dual unitarity”. In: *Phys. Rev. B* 106 (20 Nov. 2022), p. L201104.
 - [168] M. A. Rampp, S. A. Rather, and P. W. Claeys. “Geometric constructions of generalized dual-unitary circuits from biunitarity”. In: *SciPost Physics* 18.6 (June 2025).
 - [169] E. B. Mpemba and D. G. Osborne. “Cool?” In: *Phys. Educ.* 4.3 (May 1969), p. 172.
 - [170] G. Teza, J. Bechhoefer, A. Lasanta, et al. “Speedups in nonequilibrium thermal relaxation: Mpemba and related effects”. In: (2025).
 - [171] Y.-H. Ahn, H. Kang, D.-Y. Koh, et al. “Experimental verifications of Mpemba-like behaviors of clathrate hydrates”. In: *Korean J. Chem. Eng.* 33.6 (June 2016), pp. 1903–1907.
 - [172] C. Hu, J. Li, S. Huang, et al. “Conformation Directed Mpemba Effect on Polylactide Crystallization”. In: *Cryst. Growth Des.* 18.10 (2018), pp. 5757–5762.
 - [173] P. Chaddah, S. Dash, K. Kumar, et al. *Overtaking while approaching equilibrium*. 2010.
 - [174] Z. Lu and O. Raz. “Nonequilibrium thermodynamics of the Markovian Mpemba effect and its inverse”. In: *Proc. Natl. Acad. Sci. U. S. A.* 114.20 (2017), pp. 5083–5088.
 - [175] I. Klich, O. Raz, O. Hirschberg, et al. “Mpemba Index and Anomalous Relaxation”. In: *Phys. Rev. X* 9 (2 June 2019), p. 021060.
 - [176] A. Gal and O. Raz. “Precooling Strategy Allows Exponentially Faster Heating”. In: *Phys. Rev. Lett.* 124 (6 Feb. 2020), p. 060602.
 - [177] G. Teza, R. Yaacoby, and O. Raz. “Relaxation Shortcuts through Boundary Coupling”. In: *Phys. Rev. Lett.* 131 (1 July 2023), p. 017101.
 - [178] A. Kumar and J. Bechhoefer. “Exponentially faster cooling in a colloidal system”. In: *Nature* 584.7819 (Aug. 2020), pp. 64–68.
 - [179] A. Kumar, R. Chétrite, and J. Bechhoefer. “Anomalous heating in a colloidal system”. In: *Proc. Natl. Acad. Sci. U. S. A.* 119.5 (2022), e2118484119.
 - [180] I. Malhotra and H. Löwen. *Double Mpemba effect in the cooling of trapped colloids*. 2024.
 - [181] A. Lasanta, F. Vega Reyes, A. Prados, et al. “When the Hotter Cools More Quickly: Mpemba Effect in Granular Fluids”. In: *Phys. Rev. Lett.* 119 (14 Oct. 2017), p. 148001.
 - [182] A. Biswas, V. V. Prasad, O. Raz, et al. “Mpemba effect in driven granular Maxwell gases”. In: *Phys. Rev. E* 102 (1 July 2020), p. 012906.
 - [183] Z.-Y. Yang and J.-X. Hou. “Non-Markovian Mpemba effect in mean-field systems”. In: *Phys. Rev. E* 101 (5 May 2020), p. 052106.
 - [184] S. Zhang and J.-X. Hou. “Theoretical model for the Mpemba effect through the canonical first-order phase transition”. In: *Phys. Rev. E* 106 (3 Sept. 2022), p. 034131.

-
- [185] M. Jeng. “The Mpemba effect: When can hot water freeze faster than cold?” In: *American Journal of Physics* 74.6 (June 2006), pp. 514–522.
 - [186] H. C. Burridge and O. Hallstadius. “Observing the Mpemba effect with minimal bias and the value of the Mpemba effect to scientific outreach and engagement”. In: *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 476.2236 (2020), p. 20190829.
 - [187] A. Nava and M. Fabrizio. “Lindblad dissipative dynamics in the presence of phase coexistence”. In: *Phys. Rev. B* 100 (12 Sept. 2019), p. 125102.
 - [188] F. Carollo, A. Lasanta, and I. Lesanovsky. “Exponentially Accelerated Approach to Stationarity in Markovian Open Quantum Systems through the Mpemba Effect”. In: *Phys. Rev. Lett.* 127 (6 Aug. 2021), p. 060401.
 - [189] R. Bao and Z. Hou. *Accelerating relaxation in Markovian open quantum systems through quantum reset processes*. 2022.
 - [190] S. Kochsiek, F. Carollo, and I. Lesanovsky. “Accelerating the approach of dissipative quantum spin systems towards stationarity through global spin rotations”. In: *Phys. Rev. A* 106 (1 July 2022), p. 012207.
 - [191] F. Ivander, N. Anto-Sztrikacs, and D. Segal. “Hyperacceleration of quantum thermalization dynamics by bypassing long-lived coherences: An analytical treatment”. In: *Phys. Rev. E* 108 (1 July 2023), p. 014130.
 - [192] X. Wang and J. Wang. “Mpemba effects in nonequilibrium open quantum systems”. In: *Phys. Rev. Res.* 6 (3 Sept. 2024), p. 033330.
 - [193] S. Aharony Shapira, Y. Shapira, J. Markov, et al. “Inverse Mpemba Effect Demonstrated on a Single Trapped Ion Qubit”. In: *Phys. Rev. Lett.* 133 (1 July 2024), p. 010403.
 - [194] M. Moroder, O. Culhane, K. Zawadzki, et al. “Thermodynamics of the Quantum Mpemba Effect”. In: *Phys. Rev. Lett.* 133 (14 Oct. 2024), p. 140404.
 - [195] D. J. Strachan, A. Purkayastha, and S. R. Clark. “Non-Markovian Quantum Mpemba Effect”. In: *Phys. Rev. Lett.* 134 (22 June 2025), p. 220403.
 - [196] I. Medina, O. Culhane, F. C. Binder, et al. “Anomalous Discharging of Quantum Batteries: The Ergotropic Mpemba Effect”. In: *Phys. Rev. Lett.* 134 (22 June 2025), p. 220402.
 - [197] M. Xu, Z. Wei, X.-P. Jiang, et al. *Expedited thermalization dynamics in incommensurate systems*. 2025.
 - [198] P. Westhoff, S. Paeckel, and M. Moroder. *Fast and direct preparation of a genuine lattice BEC via the quantum Mpemba effect*. 2025.
 - [199] F. Ares, S. Murciano, and P. Calabrese. “Entanglement asymmetry as a probe of symmetry breaking”. In: *Nat. Commun.* 14.1 (Apr. 2023), p. 2036.
 - [200] F. Ares, P. Calabrese, and S. Murciano. “The quantum Mpemba effects”. In: *Nat. Rev. Phys.* (July 2025).
 - [201] A. Nava and R. Egger. “Pontus-Mpemba Effects”. In: *Physical Review Letters* 135.14 (Oct. 2025).
 - [202] B. Coecke, T. Fritz, and R. W. Spekkens. “A mathematical theory of resources”. In: *Inf. Comput.* 250 (2016), pp. 59–86.

-
- [203] E. Chitambar and G. Gour. “Quantum resource theories”. In: *Rev. Mod. Phys.* 91 (2 Apr. 2019), p. 025001.
 - [204] I. Marvian and R. W. Spekkens. “Modes of asymmetry: The application of harmonic analysis to symmetric quantum dynamics and quantum reference frames”. In: *Phys. Rev. A* 90 (6 Dec. 2014), p. 062110.
 - [205] E. Davies. “Generators of dynamical semigroups”. In: *J. Funct. Anal.* 34.3 (1979), pp. 421–432.
 - [206] J. Goold, M. Huber, A. Riera, et al. “The role of quantum information in thermodynamics—a topical review”. In: *J. Phys. A: Math. Theor.* 49.14 (Feb. 2016), p. 143001.
 - [207] R. Horodecki, P. Horodecki, M. Horodecki, et al. “Quantum entanglement”. In: *Rev. Mod. Phys.* 81 (2 June 2009), pp. 865–942.
 - [208] P. Faist, J. Oppenheim, and R. Renner. “Gibbs-preserving maps outperform thermal operations in the quantum regime”. In: *New Journal of Physics* 17.4 (Apr. 2015), p. 043003.
 - [209] M. Mosonyi and F. Hiai. “On the Quantum Rényi Relative Entropies and Related Capacity Formulas”. In: *IEEE Trans. Inf. Theory* 57.4 (2011), pp. 2474–2487.
 - [210] G. Gour and R. W. Spekkens. “The resource theory of quantum reference frames: manipulations and monotones”. In: *New Journal of Physics* 10.3 (2008), p. 033023.
 - [211] M. J. Donald. “Free energy and the relative entropy”. In: *J. Stat. Phys.* 49 (1987), pp. 81–87.
 - [212] F. G. S. L. Brandão, M. Horodecki, J. Oppenheim, et al. “Resource Theory of Quantum States Out of Thermal Equilibrium”. In: *Phys. Rev. Lett.* 111 (25 Dec. 2013), p. 250404.
 - [213] M. Lostaglio, D. Jennings, and T. Rudolph. “Description of quantum coherence in thermodynamic processes requires constraints beyond free energy”. In: *Nat. Comm.* 6.1 (Mar. 2015), p. 6383.
 - [214] M. Esposito, K. Lindenberg, and C. Van den Broeck. “Entropy production as correlation between system and reservoir”. In: *New J. Phys.* 12.1 (Jan. 2010), p. 013013.
 - [215] U. Seifert. “Stochastic thermodynamics, fluctuation theorems and molecular machines”. In: *Rep. Prog. Phys.* 75.12 (Nov. 2012), p. 126001.
 - [216] M. Lostaglio. “An introductory review of the resource theory approach to thermodynamics”. In: *Rep. Prog. Phys.* 82.11 (Oct. 2019), p. 114001.
 - [217] D. J. C. MacKay. *Information Theory, Inference & Learning Algorithms*. USA: Cambridge University Press, 2002.
 - [218] N. H. Y. Ng and M. P. Woods. “Resource Theory of Quantum Thermodynamics: Thermal Operations and Second Laws”. In: *Thermodynamics in the Quantum Regime: Fundamental Aspects and New Directions*. Ed. by F. Binder, L. A. Correa, C. Gogolin, et al. Cham: Springer International Publishing, 2018, pp. 625–650.
 - [219] T. Van Vu and H. Hayakawa. “Thermomajorization Mpemba Effect”. In: *Phys. Rev. Lett.* 134 (10 Mar. 2025), p. 107101.
 - [220] A. Nava and R. Egger. “Mpemba Effects in Open Nonequilibrium Quantum Systems”. In: *Phys. Rev. Lett.* 133 (13 Sept. 2024), p. 136302.

-
- [221] D. C. Alyürük, M. H. Yeşiller, V. Vedral, et al. *Thermodynamic limits of the Mpemba effect: A unified resource theory analysis*. 2025.
 - [222] G. Teza, R. Yaacoby, and O. Raz. “Eigenvalue Crossing as a Phase Transition in Relaxation Dynamics”. In: *Phys. Rev. Lett.* **130** (20 May 2023), p. 207103.
 - [223] M. Manhart. *Markov Processes*. Lecture notes for Physics 677, Rutgers University. 2016.
 - [224] A. Libál, C. Nisoli, and C. J. Olson Reichhardt. “Dynamic Control of Topological Defects in Artificial Colloidal Ice”. In: *Scientific Reports* **7** (2017).
 - [225] M. Tomamichel. *Quantum Information Processing with Finite Resources*. Springer International Publishing, 2016.
 - [226] M. Lostaglio, Á. M. Alhambra, and C. Perry. “Elementary Thermal Operations”. In: *Quantum* **2** (Feb. 2018), p. 52.
 - [227] R. Alicki. “On the detailed balance condition for non-hamiltonian systems”. In: *Reports on Mathematical Physics* **10.2** (1976), pp. 249–258.
 - [228] D. Qian, H. Wang, and J. Wang. *Intrinsic Quantum Mpemba Effect in Markovian Systems and Quantum Circuits*. 2024.
 - [229] T. Bhore, L. Su, I. Martin, et al. *Quantum Mpemba effect without global symmetries*. 2025.
 - [230] P. O’Donovan, P. Strasberg, K. Modi, et al. *Quantum master equation from the eigenstate thermalization hypothesis*. 2024.
 - [231] M. Lostaglio, K. Korzekwa, and A. Milne. “Markovian evolution of quantum coherence under symmetric dynamics”. In: *Phys. Rev. A* **96** (3 Sept. 2017), p. 032109.
 - [232] C. Rylands, K. Klobas, F. Ares, et al. “Microscopic Origin of the Quantum Mpemba Effect in Integrable Systems”. In: *Phys. Rev. Lett.* **133** (1 July 2024), p. 010401.
 - [233] S. Murciano, F. Ares, I. Klich, et al. “Entanglement asymmetry and quantum Mpemba effect in the XY spin chain”. In: *J. Stat. Mech. Theor. Exp.* **2024.1** (Jan. 2024), p. 013103.
 - [234] G. Di Giulio, X. Turkeshi, and S. Murciano. “Measurement-Induced Symmetry Restoration and Quantum Mpemba Effect”. In: *Entropy* **27.4** (Apr. 2025), p. 407.
 - [235] S. Yamashika, F. Ares, and P. Calabrese. “Entanglement asymmetry and quantum Mpemba effect in two-dimensional free-fermion systems”. In: *Phys. Rev. B* **110** (8 Aug. 2024), p. 085126.
 - [236] X. Turkeshi, P. Calabrese, and A. De Luca. “Quantum Mpemba effect in random circuits”. In: *Phys. Rev. Lett.* (June 2025).
 - [237] S. Liu, H.-K. Zhang, S. Yin, et al. “Symmetry Restoration and Quantum Mpemba Effect in Symmetric Random Circuits”. In: *Phys. Rev. Lett.* **133** (14 Oct. 2024), p. 140405.
 - [238] F. Ares, S. Murciano, P. Calabrese, et al. “Entanglement asymmetry dynamics in random quantum circuits”. In: *Physical Review Research* **7.3** (Aug. 2025).
 - [239] L. K. Joshi, J. Franke, A. Rath, et al. “Observing the Quantum Mpemba Effect in Quantum Simulations”. In: *Phys. Rev. Lett.* **133** (1 July 2024), p. 010402.
 - [240] Y. Kosmann-Schwarzbach and B. E. Schwarzbach. *The Noether Theorems: Invariance and Conservation Laws in the Twentieth Century*. Springer New York, 2011.

-
- [241] E. Noether. “Invariante Variationsprobleme”. In: *Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen, Mathematisch-Physikalische Klasse* 1918 (1918), pp. 235–257.
 - [242] V. V. Albert and L. Jiang. “Symmetries and conserved quantities in Lindblad master equations”. In: *Phys. Rev. A* 89 (2 Feb. 2014), p. 022118.
 - [243] K. Klobas, C. De Fazio, and J. P. Garrahan. “Exact pretransition effects in kinetically constrained circuits: Dynamical fluctuations in the Floquet-East model”. In: *Phys. Rev. E* 110 (2 Aug. 2024), p. L022101.
 - [244] F. Ferro, F. Ares, and P. Calabrese. “Non-equilibrium entanglement asymmetry for discrete groups: the example of the XY spin chain”. In: *J. Stat. Mech. Theor. Exp.* 2024.2 (Feb. 2024), p. 023101.
 - [245] B. Buča and T. Prosen. “A note on symmetry reductions of the Lindblad equation: transport in constrained open spin chains”. In: *New J. Phys.* 14.7 (July 2012), p. 073007.
 - [246] M. Keyl and R. F. Werner. “Optimal cloning of pure states, testing single clones”. In: *Journal of Mathematical Physics* 40.7 (July 1999), pp. 3283–3299.
 - [247] I. Marvian. “Symmetry, Asymmetry and Quantum Information, PhD thesis”. PhD thesis. <https://uwspace.uwaterloo.ca/handle/10012/7088>: University of Waterloo, 2012.
 - [248] I. Marvian and R. W. Spekkens. “Extending Noether’s theorem by quantifying the asymmetry of quantum states”. In: *Nat. Comm.* 5.1 (May 2014), p. 3821.
 - [249] A. Frigerio. “Stationary states of quantum dynamical semigroups”. In: *Communications in Mathematical Physics* 63.3 (Oct. 1978), pp. 269–276.
 - [250] H.-P. Breuer, D. Burgarth, and F. Petruccione. “Non-Markovian dynamics in a spin star system: Exact solution and approximation techniques”. In: *Phys. Rev. B* 70 (4 July 2004), p. 045323.
 - [251] I. Marvian and R. W. Spekkens. “How to quantify coherence: Distinguishing speakable and unspeakable notions”. In: *Phys. Rev. A* 94 (5 Nov. 2016), p. 052324.
 - [252] D. Janzing. “Quantum Thermodynamics with Missing Reference Frames: Decompositions of Free Energy Into Non-Increasing Components”. In: *Journal of Statistical Physics* 125.3 (Nov. 2006), pp. 761–776.
 - [253] S. D. Bartlett, T. Rudolph, and R. W. Spekkens. “Reference frames, superselection rules, and quantum information”. In: *Rev. Mod. Phys.* 79 (2 Apr. 2007), pp. 555–609.
 - [254] I. Marvian and R. W. Spekkens. “The theory of manipulations of pure state asymmetry: I. Basic tools, equivalence classes and single copy transformations”. In: *New J. Phys.* 15.3 (Mar. 2013), p. 033001.
 - [255] G. Gour, I. Marvian, and R. W. Spekkens. “Measuring the quality of a quantum reference frame: The relative entropy of frameness”. In: *Phys. Rev. A* 80 (1 July 2009), p. 012307.
 - [256] J. A. Vaccaro, F. Anselmi, H. M. Wiseman, et al. “Tradeoff between extractable mechanical work, accessible entanglement, and ability to act as a reference system, under arbitrary superselection rules”. In: *Phys. Rev. A* 77 (3 Mar. 2008), p. 032114.
 - [257] I. Marvian. “Restrictions on realizable unitary operations imposed by symmetry and locality”. In: *Nature Physics* 18.3 (2022), pp. 283–289.

-
- [258] I. Marvian. “Theory of quantum circuits with Abelian symmetries”. In: *Phys. Rev. Res.* 6 (4 Dec. 2024), p. 043292.
 - [259] D. Burgarth, G. Chiribella, V. Giovannetti, et al. “Ergodic and mixing quantum channels in finite dimensions”. In: *New J. Phys.* 15.7 (July 2013), p. 073045.
 - [260] H.-P. Breuer, E.-M. Laine, J. Piilo, et al. “Colloquium: Non-Markovian dynamics in open quantum systems”. In: *Rev. Mod. Phys.* 88 (2 Apr. 2016), p. 021002.
 - [261] R.-A. Chang, H. Shrotriya, W. W. Ho, et al. “Deep Thermalization under Charge-Conserving Quantum Dynamics”. In: *PRX Quantum* 6 (2 June 2025), p. 020343.
 - [262] K. Wold, Z. Zhu, F. Jin, et al. *Experimental Detection of Dissipative Quantum Chaos*. 2025.
 - [263] J. P. Santos, L. C. Céleri, G. T. Landi, et al. “The role of quantum coherence in non-equilibrium entropy production”. In: *Npj Quantum Inf.* 5.1 (Mar. 2019).
 - [264] S. Bravyi and A. Kitaev. “Universal quantum computation with ideal Clifford gates and noisy ancillas”. In: *Phys. Rev. A* 71 (2 Feb. 2005), p. 022316.
 - [265] S. Bravyi, G. Smith, and J. Smolin. “Trading Classical and Quantum Computational Resources”. In: *Phys. Rev. X* 6.2 (2016), p. 021043.
 - [266] Z.-W. Liu and A. Winter. “Many-Body Quantum Magic”. In: *PRX Quantum* 3.2 (May 2022).
 - [267] L. Leone, S. F. E. Oliviero, and A. Hamma. “Stabilizer Rényi Entropy”. In: *Phys. Rev. Lett.* 128 (5 Feb. 2022), p. 050402.
 - [268] L. Lumia, E. Tirrito, R. Fazio, et al. “Measurement-induced transitions beyond Gaussianity: A single particle description”. In: *Phys. Rev. Res.* 6 (2 May 2024), p. 023176.
 - [269] M. Hebenstreit, R. Jozsa, B. Kraus, et al. “All Pure Fermionic Non-Gaussian States Are Magic States for Matchgate Computations”. In: *Phys. Rev. Lett.* 123 (8 Aug. 2019), p. 080503.
 - [270] X. Lyu and K. Bu. *Displaced Fermionic Gaussian States and their Classical Simulation*. 2024.
 - [271] P. Sierant, P. Stornati, and X. Turkeshi. *Fermionic Magic Resources of Quantum Many-Body Systems*. 2025.
 - [272] X. Chen and Q. Lei. “Imaginariness Measures Induced by Relative Entropy”. In: *Reports on Mathematical Physics* 95.1 (Feb. 2025), pp. 1–10.
 - [273] J. Xu. “Quantifying the imaginarity of quantum states via Tsallis relative entropy”. In: *Physics Letters A* 528 (Dec. 2024), p. 130024.
 - [274] Y. Sun, R. Ren, Y. Wang, et al. “Analysis of the relationship between imaginarity and entanglement”. In: *Physical Review A* 111.3 (Mar. 2025).
 - [275] C. Wu and Z. Wu. “Two imaginarity monotones induced by unified (α, β) -relative entropy”. In: *Communications in Theoretical Physics* 77.9 (May 2025), p. 095101.
 - [276] A. Hickey and G. Gour. “Quantifying the imaginarity of quantum mechanics”. In: *Journal of Physics A: Mathematical and Theoretical* 51.41 (Sept. 2018), p. 414009.
 - [277] K.-D. Wu, T. V. Kondra, S. Rana, et al. “Operational Resource Theory of Imaginarity”. In: *Physical Review Letters* 126.9 (Mar. 2021).

-
- [278] K.-D. Wu, T. V. Kondra, S. Rana, et al. “Resource theory of imaginarity: Quantification and state conversion”. In: *Physical Review A* 103.3 (Mar. 2021).
 - [279] A. Belenchia, G. M. D’Ariano, and P. Perinotti. “Universality of computation in real quantum theory”. In: *EPL (Europhysics Letters)* 104.2 (Oct. 2013), p. 20006.
 - [280] N. Delfosse, P. Allard Guerin, J. Bian, et al. “Wigner Function Negativity and Contextuality in Quantum Computation on Rebits”. In: *Physical Review X* 5.2 (Apr. 2015).
 - [281] B. M. Terhal and D. P. DiVincenzo. “Classical simulation of noninteracting-fermion quantum circuits”. In: *Physical Review A* 65.3 (Mar. 2002).
 - [282] R. Jozsa and A. Miyake. “Matchgates and classical simulation of quantum circuits”. In: *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 464.2100 (July 2008), pp. 3089–3106.
 - [283] B. Dias and R. Koenig. “Classical simulation of non-Gaussian fermionic circuits”. In: *Quantum* 8 (May 2024), p. 1350.
 - [284] I. Peschel. “Calculation of reduced density matrices from correlation functions”. In: *Journal of Physics A: Mathematical and General* 36.14 (Mar. 2003), pp. L205–L208.
 - [285] S. Aaronson and D. Gottesman. “Improved simulation of stabilizer circuits”. In: *Physical Review A* 70.5 (Nov. 2004).
 - [286] V. Veitch, C. Ferrie, D. Gross, et al. “Negative quasi-probability as a resource for quantum computation”. In: *New Journal of Physics* 14.11 (Nov. 2012), p. 113011.
 - [287] D. Gross. “Hudson’s theorem for finite-dimensional quantum systems”. In: *Journal of Mathematical Physics* 47.12 (Dec. 2006).
 - [288] V. Veitch, S. A. Hamed Mousavian, D. Gottesman, et al. “The resource theory of stabilizer quantum computation”. In: *New Journal of Physics* 16.1 (Jan. 2014), p. 013009.
 - [289] X. Shi. “Erasing imaginarity: An operational method”. In: *Physical Review A* 111.5 (May 2025).
 - [290] D. A. Roberts and D. Stanford. “Diagnosing Chaos Using Four-Point Functions in Two-Dimensional Conformal Field Theory”. In: *Physical Review Letters* 115.13 (Sept. 2015).
 - [291] P. Zanardi. “Entanglement of quantum evolutions”. In: *Phys. Rev. A* 63 (4 Mar. 2001), p. 040304.
 - [292] J. Dubail. “Entanglement scaling of operators: a conformal field theory approach, with a glimpse of simulability of long-time dynamics in 1+1d”. In: *Journal of Physics A: Mathematical and Theoretical* 50.23 (May 2017), p. 234001.
 - [293] D. E. Parker, X. Cao, A. Avdoshkin, et al. “A Universal Operator Growth Hypothesis”. In: *Physical Review X* 9.4 (Oct. 2019).
 - [294] E. Rabinovici, A. Sánchez-Garrido, R. Shir, et al. *Krylov Complexity*. 2025.
 - [295] P. Nandy, A. S. Matsoukas-Roubeas, P. Martínez-Azcona, et al. “Quantum dynamics in Krylov space: Methods and applications”. In: *Physics Reports* 1125-1128 (June 2025), pp. 1–82.
 - [296] T. Prosen. “Exact Nonequilibrium Steady State of a Strongly Driven Open XXZ Chain”. In: *Physical Review Letters* 107.13 (Sept. 2011).

-
- [297] G. Pinna, O. Lunt, and C. von Keyserlingk. “Approximation theory for Green’s functions via the Lanczos algorithm”. In: *Physical Review B* 112.5 (Aug. 2025).
 - [298] M. Füllgraf, J. Wang, R. Steinigeweg, et al. “Lanczos-Pascal Approach to Correlation Functions in Chaotic Quantum Systems”. In: *Physical Review Letters* 135.25 (Dec. 2025).
 - [299] R. Haydock, V. Heine, and M. J. Kelly. “Electronic structure based on the local atomic environment for tight-binding bands”. In: *Journal of Physics C: Solid State Physics* 5.20 (Oct. 1972), pp. 2845–2858.
 - [300] V. S. Viswanath and G. Müller. *The Recursion Method: Application to Many-Body Dynamics*. Springer Berlin Heidelberg, 1994.
 - [301] J. Wang, M. Füllgraf, and J. Gemmer. “Estimate of Equilibration Times of Quantum Correlation Functions in the Thermodynamic Limit Based on Lanczos Coefficients”. In: *Physical Review Letters* 135.1 (July 2025).
 - [302] H. Mori. “A Continued-Fraction Representation of the Time-Correlation Functions”. In: *Progress of Theoretical Physics* 34.3 (Sept. 1965), pp. 399–416.
 - [303] R. Sasaki. “Towards verifications of Krylov complexity”. In: *Progress of Theoretical and Experimental Physics* (May 2024).
 - [304] B. L. Español and D. A. Wisniacki. “Assessing the saturation of Krylov complexity as a measure of chaos”. In: *Physical Review E* 107.2 (Feb. 2023).
 - [305] K. Hashimoto, K. Murata, N. Tanahashi, et al. “Krylov complexity and chaos in quantum mechanics”. In: *Journal of High Energy Physics* 2023.11 (Nov. 2023).
 - [306] F. Ballar Trigueros and C.-J. Lin. “Krylov complexity of many-body localization: Operator localization in Krylov basis”. In: *SciPost Physics* 13.2 (Aug. 2022).
 - [307] D. J. Yates and A. Mitra. “Strong and almost strong modes of Floquet spin chains in Krylov subspaces”. In: *Phys. Rev. B* 104 (19 Nov. 2021), p. 195121.
 - [308] D. J. Yates, A. G. Abanov, and A. Mitra. “Long-lived period-doubled edge modes of interacting and disorder-free Floquet spin chains”. In: *Communications Physics* 5.1 (Feb. 2022).
 - [309] P. Suchsland, R. Moessner, and P. W. Claeys. “Krylov complexity and Trotter transitions in unitary circuit dynamics”. In: *Phys. Rev. B* 111 (1 Jan. 2025), p. 014309.
 - [310] A. A. Nizami and A. W. Shrestha. “Krylov construction and complexity for driven quantum systems”. In: *Physical Review E* 108.5 (Nov. 2023).
 - [311] G. Teschl. *Jacobi Operators and Completely Integrable Nonlinear Lattices*. American Mathematical Society, Oct. 1999.
 - [312] S. E. Tapias Arze, P. W. Claeys, I. Pérez Castillo, et al. “Out-of-equilibrium phase transitions induced by Floquet resonances in a periodically quench-driven XY spin chain”. In: *SciPost Physics Core* 3.1 (July 2020).
 - [313] P. M. Schindler and M. Bukov. “Geometric Floquet Theory”. In: *Physical Review X* 15.3 (Aug. 2025).
 - [314] J. H. Shirley. “Solution of the Schrödinger Equation with a Hamiltonian Periodic in Time”. In: *Phys. Rev.* 138 (4B May 1965), B979–B987.

-
- [315] H. Sambe. “Steady States and Quasienergies of a Quantum-Mechanical System in an Oscillating Field”. In: *Phys. Rev. A* 7 (6 June 1973), pp. 2203–2213.
 - [316] M. Grifoni and P. Hänggi. “Driven quantum tunneling”. In: *Physics Reports* 304.5 (1998), pp. 229–354.
 - [317] W. Magnus. “On the exponential solution of differential equations for a linear operator”. In: *Communications on Pure and Applied Mathematics* 7.4 (Nov. 1954), pp. 649–673.
 - [318] S. Blanes, F. Casas, J. Oteo, et al. “The Magnus expansion and some of its applications”. In: *Physics Reports* 470.5-6 (Jan. 2009), pp. 151–238.
 - [319] S. Rahav, I. Gilary, and S. Fishman. “Effective Hamiltonians for periodically driven systems”. In: *Physical Review A* 68.1 (July 2003).
 - [320] A. Eckardt and E. Anisimovas. “High-frequency approximation for periodically driven quantum systems from a Floquet-space perspective”. In: *New Journal of Physics* 17.9 (Sept. 2015), p. 093039.
 - [321] N. Goldman and J. Dalibard. “Periodically Driven Quantum Systems: Effective Hamiltonians and Engineered Gauge Fields”. In: *Physical Review X* 4.3 (Aug. 2014).
 - [322] T. Mikami, S. Kitamura, K. Yasuda, et al. “Brillouin-Wigner theory for high-frequency expansion in periodically driven systems: Application to Floquet topological insulators”. In: *Physical Review B* 93.14 (Apr. 2016).
 - [323] F. Casas, J. A. Oteo, and J. Ros. “Floquet theory: exponential perturbative treatment”. In: *Journal of Physics A: Mathematical and General* 34.16 (Apr. 2001), pp. 3379–3388.
 - [324] Y. Braver and E. Anisimovas. “Extended degenerate perturbation theory for the Floquet-Hilbert space”. In: *Physical Review A* 110.2 (Aug. 2024).
 - [325] G. D. Dinc, A. Eckardt, and A. Schnell. “Effective Floquet Lindblad generators from spectral unwinding”. In: *Physical Review A* 111.6 (June 2025).
 - [326] H. W. Kuhn. “The Hungarian method for the assignment problem”. In: *Naval Research Logistics Quarterly* 2.1-2 (Mar. 1955), pp. 83–97.
 - [327] J. Munkres. “Algorithms for the Assignment and Transportation Problems”. In: *Journal of the Society for Industrial and Applied Mathematics* 5.1 (Mar. 1957), pp. 32–38.
 - [328] M. B. Hastings and X.-G. Wen. “Quasiadiabatic continuation of quantum states: The stability of topological ground-state degeneracy and emergent gauge invariance”. In: *Physical Review B* 72.4 (July 2005).
 - [329] D. Sels and A. Polkovnikov. “Minimizing irreversible losses in quantum systems by local counterdiabatic driving”. In: *Proceedings of the National Academy of Sciences* 114.20 (May 2017).
 - [330] B. Bertini, P. Kos, and T. Prosen. “Operator entanglement in local quantum circuits I: Chaotic dual-unitary circuits”. In: *SciPost Physics* 8.4 (2020), p. 067.
 - [331] D. Gross, V. Nesme, H. Vogts, et al. “Index Theory of One Dimensional Quantum Walks and Cellular Automata”. In: *Communications in Mathematical Physics* 310.2 (Jan. 2012), pp. 419–454.
 - [332] H. C. Po, L. Fidkowski, T. Morimoto, et al. “Chiral Floquet Phases of Many-Body Localized Bosons”. In: *Physical Review X* 6.4 (Dec. 2016).

-
- [333] K. Klobas, C. Rylands, and B. Bertini. “Translation symmetry restoration under random unitary dynamics”. In: *Phys. Rev. B* 111 (14 Apr. 2025), p. L140304.
 - [334] D. Gonzalez-Cuadra, E. Tirrito, F. Petiziol, et al. “Digital Quantum Simulation of a (1+1)D SU(2) Lattice Gauge Theory with Ion Qudits”. In: *PRX Quantum* 5.4 (2024), p. 040309.
 - [335] L. Piroli, B. Bertini, J. I. Cirac, et al. “Exact dynamics in dual-unitary quantum circuits”. In: *Phys. Rev. B* 101 (9 Mar. 2020), p. 094304.
 - [336] M. Mazzoni, L. Capizzi, and L. Piroli. “Breaking global symmetries with locality-preserving operations”. In: *arXiv preprint arXiv:2508.15892* (2025).
 - [337] M. Howard and E. Campbell. “Application of a Resource Theory for Magic States to Fault-Tolerant Quantum Computing”. In: *Physical Review Letters* 118.9 (Mar. 2017).
 - [338] R. Rubboli, R. Takagi, and M. Tomamichel. “Mixed-state additivity properties of magic monotones based on quantum relative entropies for single-qubit states and beyond”. In: *Quantum* 8 (Oct. 2024), p. 1492.
 - [339] T. Baumgratz, M. Cramer, and M. B. Plenio. “Quantifying Coherence”. In: *Physical Review Letters* 113.14 (Sept. 2014).
 - [340] X. Turkeshi, E. Tirrito, and P. Sierant. “Magic spreading in random quantum circuits”. In: *Nature Communications* 16.1 (Mar. 2025), p. 2575.
 - [341] P. S. Tarabunga, M. Frau, T. Haug, et al. *A nonstabilizerness monotone from stabilizerness asymmetry*. 2024.
 - [342] G. H. Low and I. L. Chuang. “Optimal Hamiltonian Simulation by Quantum Signal Processing”. In: *Phys. Rev. Lett.* 118 (1 Jan. 2017), p. 010501.
 - [343] Y. Kikuchi, C. Mc Keever, L. Coopmans, et al. “Realization of quantum signal processing on a noisy quantum computer”. In: *npj Quantum Information* 9.1 (Sept. 2023).
 - [344] T. Holden-Dye, L. Masanes, and A. Pal. “Fundamental charges for dual-unitary circuits”. In: *Quantum* 9 (2025), p. 1615.
 - [345] M. Polettini and M. Esposito. “Nonconvexity of the relative entropy for Markov dynamics: A Fisher information approach”. In: *Phys. Rev. E* 88 (1 July 2013), p. 012112.
 - [346] Y. Cao, J. Lu, and Y. Lu. “Gradient flow structure and exponential decay of the sandwiched Rényi divergence for primitive Lindblad equations with GNS-detailed balance”. In: *Journal of Mathematical Physics* 60.5 (May 2019), p. 052202.