

Numerical methods for many-body quantum dynamics

Thesis by
Christopher David White

In Partial Fulfillment of the Requirements for the
Degree of
Doctor of Philosophy

The logo for the California Institute of Technology (Caltech), featuring the word "Caltech" in a bold, orange, sans-serif font.

CALIFORNIA INSTITUTE OF TECHNOLOGY
Pasadena, California

2019
Defended 16 May 2019

© 2019

Christopher David White
ORCID: 0000-0002-8372-2492

Some rights reserved. This thesis is distributed under a Creative Commons
Attribution NonCommercial-ShareAlike License.

ACKNOWLEDGEMENTS

cum caritate et gratia magna parochiae et populo Sancti Thomas Apostoli Holivudensis. “congregavit nos in unum Christi amor; exsulemus, et in ipso jucundemur.”

ABSTRACT

This thesis describes two studies of the dynamics of many-body quantum systems with extensive numerical support.

In Part I we first give a new algorithm for simulating the dynamics of one-dimensional systems that thermalize (that is, come to local thermal equilibrium). The core of this algorithm is a new truncation for matrix product operators, which reproduces local properties faithfully without reproducing non-local properties (e.g. the information required for OTOCs). To the extent that the dynamics depends only on local operators, timesteps interleaved with this truncation will reproduce that dynamics.

We then apply this algorithm to Floquet systems: first to clean, non-integrable systems with a high-frequency drive, where we find that the system is well-described by a natural diffusive phenomenology; and then to disordered systems with low-frequency drive, which display diffusion—not subdiffusion—at appreciable disorder strengths.

In Part II, we study the utility of many-body localization as a medium for a thermodynamic engine. We first construct a small (“mesoscale”) engine that gives work at high efficiency in the adiabatic limit, and show that thanks to the slow spread of information in many body localized systems, these mesoscale engines can be chained together without specially engineered insulation. Our construction takes advantage of precisely the fact that MBL systems do *not* thermalize. We then show that these engines still have high efficiency when run at finite speed, and we compare to competitor engines.

PUBLISHED CONTENT AND CONTRIBUTIONS

This thesis contains material from:

White, C.D., Zaletel, M., Mong, R.S.K., and Refael, G. “Quantum Dynamics of Thermalizing Systems.” *Phys. Rev. B* 97, 035127 (2018). DOI: 10.1103/physrevb.97.035127.

CDW constructed the algorithm used, wrote the algorithm-specific code, constructed the benchmarking protocol, carried out the simulations, and performed the analysis.

Ye, B., Machado, F., White, C.D., Mong, R.S.K., and Yao, N. “Emergent hydrodynamics in Floquet quantum systems”. [arXiv:1902.01859](https://arxiv.org/abs/1902.01859).

CDW participated in the conception of the project and the analysis of the results (in particular the framing in terms of hydrodynamics). In addition, the project used CDW’s algorithm-specific code.

Yunger Halpern, N., White, C.D., Gopalakrishnan, S., and Refael, G. “Quantum engine based on many-body localization”. *Phys. Rev. B*. 99, 024203. DOI: 10.1103/physrevb.99.024203

CDW participated in the conception of the project, carried out all of the numerical calculations and many of the analytical calculations, and participated in constructing the physical arguments.

TABLE OF CONTENTS

Acknowledgements	iii
Abstract	iv
Published Content and Contributions	v
Table of Contents	vi
Introduction	ix
Chapter I: Background: MBL and MPS	1
1.1 Many body localization	1
1.1.1 Many-body localization in perturbation theory	1
1.1.2 Long-time memory in MBL systems	5
1.1.3 Level spacing	7
1.1.4 Local dynamics at finite time	10
1.2 Matrix product states	10
1.2.1 Graphical notation for tensors and tensor contractions	14
1.2.2 Expectation values and correlations in MPSs	15
1.2.3 Canonical forms and orthogonality centers	16
1.3 Dynamics and matrix product states	20
1.4 Matrix product operators	21
1.4.1 MPO representation of an exponential interaction	22
1.4.2 MPO representation of a power-law Hamiltonian	23
1.4.3 MPDOs	24
1.4.4 Matrix product density operator representations of Gibbs states	26
1.5 Some essentials of quantum thermodynamics	26
I Thermalizing systems and matrix product density operators	28
Chapter II: Simulating quantum dynamics of thermalizing systems	29
2.1 Background and intuition	30
2.1.1 Background: matrix product state methods	30
2.1.2 Intuition: thermalization and computation	33
2.2 Method: truncation of MPDOs	36
2.2.1 Setting, notation, and tools	37
2.2.2 Re-writing the MPDO to expose properties whose preservation we guarantee	39
2.2.3 Modifying the MPDO	41
2.2.4 Preservation of l -site operators	44
2.3 Results	46
2.3.1 Hamiltonian	46

2.3.2	Application: pure-state evolution	47
2.3.3	Application: mixed-state evolution (near equilibrium)	52
2.3.4	Application: mixed-state evolution (far from equilibrium)	52
2.4	Conclusion	55
Chapter III: Heating, prethermalization, and hydrodynamics in a high-frequency Floquet system		59
3.1	Model and Floquet phenomenology	60
3.2	Benchmarking DMT	65
3.2.1	Benchmarking Krylov subspace dynamics	65
3.2.2	DMT and Krylov	66
3.2.3	Entropy in DMT	69
3.2.4	Effect of Trotter step size on DMT numerics	69
3.3	Analysis and results	71
3.3.1	Approach to Gibbs Ensemble	73
3.3.2	Heating	77
3.3.3	Extracting diffusion coefficients of a spatially uniform, static Hamiltonian	78
3.3.4	Classical diffusion equation	82
3.3.5	Hydrodynamics in large spin chains	86
3.4	Discussion	86
Chapter IV: Hydrodynamics in a disordered low-frequency Floquet system		88
4.1	Model	89
4.2	Method	90
4.2.1	Convergence testing	91
4.2.2	Comparison with exact simulations	93
4.2.3	Quantity of interest	95
4.2.4	Disorder averaging	96
4.3	Results	97
4.4	Discussion	99

II MBL-mobile: Many-body-localized engine 101

Chapter V: MBL and thermodynamic engines		102
Chapter VI: The MBL-Mobile in its adiabatic limit		105
6.1	Qubit toy model for the mesoscale engine	105
6.2	Set-up for the mesoscale MBL engine	107
6.3	Notation and definitions:	110
6.4	Quantitative analysis of the adiabatic mesoscale engine: some easy limits	111
6.5	Quantitative analysis of the adiabatic mesoscale engine: a more detailed calculation	113
6.5.1	Partial-swap model of thermalization	113
6.5.2	Average heat $\langle Q_2 \rangle$ absorbed during stroke 2	114
6.5.3	Average heat $\langle Q_4 \rangle$ absorbed during stroke 4	116
6.5.4	Average per-cycle power $\langle W_{\text{tot}} \rangle$	118

6.5.5	Efficiency η_{MBL} in the adiabatic approximation	120
6.6	MBL engine in the thermodynamic limit	121
6.7	Numerical simulations	125
6.7.1	Hamiltonian	125
6.7.2	Scaling factor	126
6.7.3	Representing states and Hamiltonians	129
6.7.4	Strokes 1 and 3: tuning	129
6.7.5	Stroke 2: Thermalization with the cold bath	130
6.7.6	Results	131
Chapter VII:	The MBL-Mobile at finite speed	133
7.1	Diabatic corrections	133
7.1.1	Fractional-Landau-Zener transitions	134
7.1.2	Landau-Zener transitions	135
7.1.3	APT transitions	136
7.2	Numerical simulations	137
7.2.1	Simulating finite-time tuning	138
7.2.2	Enlightenment from numerical methods	139
7.3	Precluding communication between subengines	140
7.4	Lower bound on the cycle time from cold thermalization	141
Chapter VIII:	Racing the MBL-mobile	144
8.1	Comparisons with competitor Otto engines	144
8.1.1	Comparison with bandwidth engine	144
8.1.2	Comparison with MBL engine tuned between same-strength disorder realizations	145
8.1.3	Quantum-dot engine	147
8.1.4	Anderson-localized engine	147
8.2	Order-of-magnitude estimates	148
8.3	Outlook	149
Bibliography	152

INTRODUCTION

Numerical methods for many-body quantum systems go back to the earliest days of electronic computers. The physicists of Los Alamos during World War II relied heavily on numerical simulations; the first calculations were done by human “computers” or with mechanical calculators,¹ but soon after the war they began using ENIAC, the first programmable electronic computer [84].

We need simulations because the phenomena we study are complex. This was certainly the case for the pioneers of computational physics, simulating weapons physics at Los Alamos, and it is the case for us now. We may have a microscopic Hamiltonian, but analytically solving the Schrödinger equation with that Hamiltonian is out of the question, and tightly-controlled approximations are frequently unavailable. So we are forced to rely on heuristic approximations justified by physical arguments. But these approximations need to be checked: do they describe the model’s behaviour well? At the very least, do they describe some essential phenomenon at work in the system? Numerical simulations of a model catch hidden assumptions and subtle mistakes in your physical arguments—and they suggest further approximations to make and phenomena to investigate. (The work of Part II of this thesis proceeded in exactly this fashion, as an ongoing three-part conversation between physical argument, careful analytical work, and numerical simulation.)

But numerical methods are not simply devices for checking your intuitions and arguments. Developing a *new* numerical method and teasing out a characterization of where it succeeds, where it fails, and why it succeeds or fails provides important insight into the physics of the systems that method simulates

Take the history of the **density matrix renormalization group (DMRG)** and the **matrix product state (MPS)**.² Steve White developed the method in 1993 [226, 227] as an extension of Wilson’s numerical renormalization group (NRG) [230], and it quickly saw wide application [82, 100, 127, 180], not only for computing ground states but also for simulating low-energy dynamics. Meanwhile, in 1995, Ostlund and Rommer [248] showed that “infinite-system” DMRG could be formulated in terms of matrix product states, which had existed for some time, e.g. as a mechanism for writing down the ground

¹The Los Alamos physicists required numerical stability from their algorithms because dust caused frequent errors.

²This paragraph draws heavily on the review [187] of Schollwöck.

state of the AKLT Hamiltonian [9, 10, 59], and in 1999 Takasaki et al. [203] showed that “finite-system” DMRG is precisely a variational procedure, which iteratively minimizes the energy over the manifold of matrix product states.

The notion of truncation was crucial to the practical developments. Matrix product states offer not only a compact representation of low-entanglement states, but also an obvious procedure for “lossy compression”—for optimally approximating one matrix product state by another that requires less memory and computation time to store and use.

The growing understanding that ground states of gapped Hamiltonians have efficient matrix product state representations (in particular the work of Hastings, e.g. [86]) gave physicists the confidence to use DMRG as, essentially, a black-box method. Since most translation-invariant local Hamiltonians are gapped (see [131] for a recent proof) and ground states of gapped Hamiltonians have efficient matrix product representations, one expects DMRG to reliably find the ground state of an arbitrarily-chosen Hamiltonian. And DMRG is generally **self-diagnosing**: even if it does fail, you can tell by checking convergence in the simulation parameters, without having to compare to an exact result, or argue without empirical evidence that it should give the right answer.³ Even if you are not a specialist, you can choose from a number of carefully-optimized, well-engineered MPS implementations, e.g. ITensor [1] or TeNPy3 [87], and expect that either DMRG will work or you will see clear evidence that it fails. The long list of papers using ITensor (<http://itensor.org/docs.CGI?page=papers>) testifies to the impact of this development.

But that understanding also led to a number of analytical results, notably Xie Chen and friends’ classification of bosonic symmetry protected topological phases (SPTs) [36]; cf Pollmann and friends’ [169, 170]. [170] is particularly close to the spirit of this thesis: in it, Pollmann and his collaborators ask how symmetry operations act on the vector space in which the entanglement spectrum of an MPS in **mixed canonical form** lives. They find that the decomposition of that vector space into irreducible projective representations of the symmetry group strongly constrains the entanglement spectrum.

³A few of DMRG’s failure modes, notably difficulty capturing rare long-range correlations as in a random-singlet ground state, are not self-diagnosing. See [184] for some of this problem and the performance of an algorithm RRG that performs better for these states; for more examples of failures in disordered systems see *TODO schmitteckert Brenden*.

And DMRG also gave rise to a variety of numerical methods for time evolution, notably TEBD [213, 214, 228]. (I discuss TEBD in some detail in Sec. 1.3, and give an opinionated discussion of a variety of MPS-based time evolution methods in Sec.2.1.1.) These numerical methods in turn shed light on one of the fundamental questions of computational physics—what physical processes can one efficiently simulate on a classical computer?—by carving out a set of quantum problems—namely those that generate little entanglement—that *can* be efficiently simulated. This question in turn is intimately related to the **extended Church-Turing thesis**. Other work along these lines includes the Gottesman-Knill theorem [73], which shows that computations involving only Clifford gates are efficiently simulable on classical computers.

A feedback loop between development and characterization of numerical methods, on the one hand, and analytical work of varying degrees of rigor on the other, is clearly a fruitful way to do research: it leads to both practical numerical methods and to a deeper analytical understanding of the phenomena we started out trying to simulate.

Until recently, this feedback loop had been applied in only the most limited of ways to questions of thermalization, disorder, and localization. Generic (viz., non-integrable) systems are expected to come to thermal equilibrium in finite time; we call these systems **thermalizing**. Thermalization is understood in terms of the **eigenstate thermalization hypothesis (ETH)**: because the time to approach a state that is effectively random except to the extent that it has a certain energy expectation value—that is, a Gibbs ensemble—is given by the exponentially long dephasing time between near-in-energy eigenstates, *each eigenstate* must have the properties of a Gibbs state at that eigenstate’s eigenenergy [44, 54, 182, 183, 200]. Interestingly, the ETH is not enough. Some systems can satisfy the ETH but have very slow thermalization times, either because they have only weak integrability-breaking terms, or because they give rise to emergent local quasi-conserved quantities [20, 113, 136]; when I say “thermalizing” I implicitly mean “quickly thermalizing”, for some notion of “quickly” determined by the other timescales of the problem at hand and by the resources (memory and computation time) available for simulations. And not all “generic” systems thermalize at all: sufficiently strongly disordered systems undergo **many-body localization (MBL)**, in the Hamiltonian commutes with local conserved quantities, essentially dressed onsite occupations. Near the transition between the ETH and MBL phases, a regime displaying rich Griffiths physics, e.g. **subdiffusion**, is expected to occur.

Because thermalizing systems generate entanglement quickly, TEBD is of limited utility, and numerical studies were for some time restricted to exact diagonalization (e.g. [115]), with occasional forays into special-purpose MPS or matrix-product operator (MPO) calculations (e.g. [113]). But in the last few years, three methods—TDVP [133], CTWA [233, 234], and my own **density matrix truncation**, or **DMT**—have **density matrix truncation** come into use; they promise to shed light on the physics of thermalizing systems.

This thesis sits in the tradition I have sketched of using both sides of this feedback loop to advance the discipline.

I present necessary background in Chapter 1. The later chapters assume some basic knowledge of many-body localization, matrix product states, and algorithms for time evolution. Chapter 1 communicates that knowledge. I aim more for intuition than rigor. Nothing in 1 is original; most of the contents are ambient cultural knowledge, so I include few citations.

Part I describes DMT and the physics we can use it to study. DMT is a new describe a new algorithm for approximating matrix product representations of density matrices that takes advantage of our understanding of thermalization. It succeeds at simulating previously unsimulable systems, and its convergence properties already prompt interesting insights into those systems.

In Ch. 2, we argue that truncating MPDOs in such a way that the expectation values of the conserved quantities of some thermalizing Hamiltonian do not change should not change the long-time dynamics of that Hamiltonian. We construct a truncation procedure that does this, and find that it does indeed capture the long-time dynamics of a fruit-fly thermalizing Hamiltonian. In Ch. 3 we apply DMT to a high-frequency Floquet system. We find that it accurately captures not only the system’s hydrodynamics but also the Floquet heating process; the degree to which it does this illuminates the details of that process. We also find that the long-time prethermal dynamics of the system are well-described by a phenomenological diffusion equation In Ch. 4, we apply DMT to a low-frequency disordered Floquet system, and find that even for substantial disorder its hydrodynamics are diffusive. We also comment on some subtleties in evaluating DMT’s accuracy, and the degree to which it is self-diagnosing.

In Part II I turn to a very different work: on turning MBL systems to thermodynamic advantage. After briefly contextualizing the problem in Ch. 5, we construct a thermodynamic engine from an MBL system and analyze it in a

relatively straightforward adiabatic limit (Ch. 6); we then analyze finite-speed effects (Ch. 7), and compare it to competitor engines (Ch. 8).

At first sight the material of Part II does not fit so well into the tradition of using methods development to comment on physics, and vice versa. But this impression understates the degree to which the numerical simulations shaped our numerical and even conceptual work, and it also understates the degree to which the *limits* of the numerical methods we used give us insight into the physics we use and the robustness of our protocols. I briefly comment on this point in Sec. 7.2

Chapter 1

BACKGROUND: MBL AND MPS

1.1 Many body localization

1.1.1 Many-body localization in perturbation theory

Consider a random paramagnet—a system with Hamiltonian

$$H_{\text{rpara}} = \sum h_j^z \sigma_j^z, \quad (1.1)$$

where the h_j^z are drawn independently from the uniform distribution on $[-h, h]$. A system governed by this state has “perfect memory” in the sense that if it starts out in some product of σ_j^z eigenstates, say

$$|0\text{xCOFFEE}\rangle = |\uparrow\uparrow\downarrow\downarrow\downarrow\downarrow\downarrow\downarrow\uparrow\uparrow\uparrow\uparrow\uparrow\uparrow\uparrow\uparrow\downarrow\uparrow\uparrow\downarrow\rangle, \quad (1.2)$$

it will remain in that state forever: $e^{-iHt} |0\text{xCOFFEE}\rangle = |0\text{xCOFFEE}\rangle$, up to a phase, for all t . In terms of operators

$$\frac{d}{dt} \langle \sigma_j^z \rangle = -i \langle [H_{\text{rpara}}, \sigma_j^z] \rangle = 0. \quad (1.3)$$

Indeed $[H_{\text{rpara}}, \sigma_j^z] = 0$: the σ_j^z are conserved quantities of the Hamiltonian H_{rpara} .

Now suppose we change the Hamiltonian. Suppose in particular that we add a nearest neighbor interaction:

$$\begin{aligned} H_{\text{RFH}} &= H_{\text{rpara}} + t \sum_j [\sigma_j^x \sigma_{j+1}^x + \sigma_j^y \sigma_{j+1}^y] + \Delta \sum_j \sigma_j^z \sigma_{j+1}^z \\ &= t \sum_j [\sigma_j^x \sigma_{j+1}^x + \sigma_j^y \sigma_{j+1}^y] + \Delta \sum_j \sigma_j^z \sigma_{j+1}^z + \sum_j h_j^z \sigma_j^z, \end{aligned} \quad (1.4)$$

h_j^z still random in $[-h, h]$. (I choose this Hamiltonian with malice aforethought. Since the work of Pal and Huse [160] this Hamiltonian, called the random-field Heisenberg Hamiltonian, has served as the fruit-fly model for the study of many body localization: see inter alia multa [246].) Then

$$[H_{\text{RFH}}, \sigma_j^z] \neq 0, \quad (1.5)$$

i.e. the σ_j^z are no longer conserved quantities. Disaster!

But if h is large, the σ_j^z are *almost* conserved. First, take $\Delta = 0$ and add $h \neq 0$. Then the Hamiltonian is

$$H_{\&} = t \sum_j [\sigma_j^x \sigma_{j+1}^x + \sigma_j^y \sigma_{j+1}^y] + \sum_j h_j^z \sigma_j^z . \quad (1.6)$$

On Jordan-Wigner transformation this is the single-particle Hamiltonian

$$H_{\&} = 2t \sum_j [c_j^\dagger c_{j+1} + \text{h.c.}] + \sum_j h_j^z n_j . \quad (1.7)$$

I'm going to take as read the fact that this Hamiltonian displays **Anderson localization**, so its single-particle eigenstates are localized with some localization length, call it ζ .

To put that statement about localized single-particle eigenstates in a more useful form, write \tilde{n}_j for the occupation number operator of the eigenstate centered at site j , and write $\tau_{\&j}^z = 2\tilde{n}_j$. (Similarly Jordan-Wigner transform $\tilde{c}^\dagger \sim \tau^+$.) $\tau_{\&j}^z$ is called an (Anderson) **l-bit**. Then

$$H_{\&} = \sum_j h'_j \tau_{\&j}^z , \quad (1.8)$$

where the h'_j are the eigenenergies of the single-particle eigenstates. There is some unitary $U_{\&}$ (namely the unitary diagonalizing $H_{\&}$) such that

$$U_{\&}^\dagger \tau_{\&j}^z U_{\&} = \sigma^z \quad (1.9)$$

and $U_{\&}^\dagger \tau_{\&j}^+ U_{\&} = \sigma^+$, etc. Since the single-particle eigenstates are localized with localization length ζ ,

$$\tau_{\&j}^z \sim \sigma_j^z + \sum_{kl} R_{kl}^j e^{-(|k-j|+|j-l|)/\zeta} \sigma_k^+ \sigma_l^- \quad (1.10)$$

for some appropriate random R_{kl}^j .

How do interactions change this picture? Add an interaction term $V = \sum_j \sigma_j^z \sigma_{j+1}^z$, for

$$H_{\text{RFH}} = H_{\&} + V = \sum_j \vec{\sigma}_j \cdot \vec{\sigma}_{j+1} + \sum_j h_j \sigma_j^z . \quad (1.11)$$

We would hope to treat this as a perturbation. We can only do so if

$$\left| \frac{\langle \alpha | V | \alpha \rangle}{D_\alpha - D_\beta} \right| \ll 1. \quad (1.12)$$

To estimate this, suppose two states $|\alpha\rangle, |\beta\rangle$ differ on a subsystem of diameter l : that is,

$$|\alpha\rangle\langle\beta| = \tau_{\&a_1}^+ \cdots \tau_{\&a_n}^+ \tau_{\&b_1}^- \cdots \tau_{\&b_m}^-, \quad (1.13)$$

where the sites $a_1, \dots, a_n, b_1, \dots, b_m$ are all in a subsystem l . Then

$$\begin{aligned} \langle \beta | \sigma_j^z \sigma_{j+1}^z | \alpha \rangle &= \text{tr} [\sigma_j^z \sigma_{j+1}^z | \alpha \rangle \langle \beta |] \\ &= \text{tr} [\sigma_j^z \sigma_{j+1}^z \tau_{\&a_1}^+ \cdots \tau_{\&a_n}^+ \tau_{\&b_1}^- \cdots \tau_{\&b_m}^-] \\ &= \text{tr} [(U_{\&}^\dagger \sigma_j^z \sigma_{j+1}^z U_{\&}) \sigma_{a_1}^+ \cdots \sigma_{a_n}^+ \sigma_{b_1}^- \cdots \sigma_{b_m}^-]. \end{aligned} \quad (1.14)$$

What is the operator content of this strange construction $U^\dagger \sigma_j^z \sigma_{j+1}^z U$? We know that conjugation by U maps

$$U_{\&} \sigma_j^z U_{\&}^\dagger \sim \frac{1}{\zeta} \sigma_j^z + \sum_{kl} R_{kl}^j e^{-(|k-j|+|j-l|)/\zeta} \sigma_k^+ \sigma_l^- \quad (1.15)$$

with $R_{kl}^j \sim \zeta^{-1}$ random, so it is reasonable to suppose that $U^\dagger \sigma_j^z \sigma_{j+1}^z U$ is

$$U^\dagger \sigma_j^z \sigma_{j+1}^z U = \sum_{\vec{\nu}} R_{\vec{\nu}}^j 2^{-2\ell[\vec{\nu}]} e^{-\ell[\vec{\nu}]/\zeta} \sigma^{\vec{\nu}}, \quad (1.16)$$

where $R_{\vec{\nu}}^j \sim \frac{1}{\zeta}$, $\vec{\nu} = \nu_1 \nu_2 \nu_3 \dots \nu_l$, $\nu_j \in 0 \dots 3$ label Pauli strings $\sigma^{\vec{\nu}}$, we take each $\vec{\nu}$ to have at least one $\nu_j \neq 0$, and $\ell[\vec{\nu}]$ is the natural notion of the diameter of the Pauli string $\sigma^{\vec{\nu}}$

$$\ell[\vec{\nu}] \equiv \max\{j : \nu_j \neq 0\} - \min\{j : \nu_j \neq 0\}. \quad (1.17)$$

The factor of $2^{-2\ell[\vec{\nu}]}$ ensures normalization. Then

$$\begin{aligned} \langle \beta | \sigma_j^z \sigma_{j+1}^z | \alpha \rangle &= \text{tr} [(U_{\&}^\dagger \sigma_j^z \sigma_{j+1}^z U_{\&}) \sigma_{a_1}^+ \cdots \sigma_{a_n}^+ \sigma_{b_1}^- \cdots \sigma_{b_m}^-] \\ &\sim 2^{-l/2} e^{-l/\zeta}. \end{aligned} \quad (1.18)$$

The energy difference is typically not less than the spacing of levels on l sites:

$$D_\alpha - D_\beta \gtrsim h l 2^{-l}. \quad (1.19)$$

So our condition for V to act perturbatively is that for most α, β

$$1 \gg \left| \frac{\langle \alpha | V | \beta \rangle}{D_\alpha - D_\beta} \right| \sim e^{l(\ln 2 - \zeta^{-1})} \quad (1.20)$$

or equivalently that the bare model $H_\&$ have localization length

$$\zeta^{-1} > \ln 2 . \quad (1.21)$$

(This is a heuristic sketch of certain elements of de Roeck et al.’s “avalanche picture”; see also [5, 49, 145, 171, 206, 207]. The chief element I have not treated is the possibility—near certainty, in fact—that some regions will be resonant in the sense that the condition (1.20) is not satisfied; de Roeck and friends give an RG-like scheme for determining whether these “resonant regions” destabilize the Anderson-localized system.)

If you carry this diagonalization out in full you will find a unitary U such that $\tau_j^z = U \tau_{\&j}^z U^\dagger = U U_\& \sigma_j^z U_\&^\dagger U_\&^\dagger$ is a conserved quantity of H_{RFH} :

$$[H_{\text{RFH}}, \tau_j^z] = 0 . \quad (1.22)$$

Based on the condition (1.20) and the expectation that higher-order terms will scale as

$$[n\text{th-order term}] \sim \left(e^{l(\ln 2 - \zeta^{-1})} \right)^n , \quad (1.23)$$

one expects these l-bits to have the form localization length

$$\tau_j^z \sim U^\dagger U_\&^\dagger \sigma_j^z U_\& U \sim \frac{1}{\xi} \sigma_j^z + \sum_{kl} R_{kl}^j e^{-(|k-j|+|j-l|)/\xi} \sigma_k^+ \sigma_l^- \quad (1.24)$$

where

$$\xi = (\ln 2 - \zeta^{-1})^{-1} . \quad (1.25)$$

(Again this is morally consistent with the results of the avalanche picture, but they run this quantity through an RG scheme.) We have therefore constructed extensively many local integrals of motion. These are l-bits in the usual sense [34, 92].

A couple of notes about the condition (1.21) on the Anderson localization length are in order. The right hand side of that condition, which ultimately comes from an estimate of the level spacing, is actually an entropy density. This has consequences:




Figure 1.1: **Left:** Notional density of states (hence entropy density) for a disordered many-body system. **Right:** Notional phase diagram, per (1.21). Compare the phase diagram found in [146] for the isotropic random-field Heisenberg model.

1. Systems with larger onsite Hilbert spaces, from the prosaic (spin-3/2 sites) to the exotic (SYK puddles), will require larger disorder to localize (if they localize at all: constructing this Anderson route to MBL will require creativity).
2. Since for real systems the entropy density is a function of energy density, the condition (1.21) (as well as the more sophisticated treatments of de Roeck et al.) naïvely predict a mobility edge. See Fig. 1.1. But de Roeck et al. argue in [48] that a mobility edge is impossible, in the sense that it will destabilize localized states at lower energies. Per this argument, interacting systems must be fully MBL (no mobility edge) or not MBL at all. On the other hand, numerics (e.g. [146]) consistently show a mobility edge consistent with the naïve prediction of (1.21)). My own belief is that that the mobility edge exists, but this will require a great deal of work to sort out.

In addition, the full avalanche picture predicts that MBL cannot exist in $d \geq 2$ dimensions. You can already glimpse how this comes about: if you construct a perturbative scaling in some length l based on (1.20), the level spacing shrinks exponentially in l^d , while the matrix element only shrinks exponentially in l . So in $d \geq 2$, couplings are always resonant at sufficiently long length scales.

1.1.2 Long-time memory in MBL systems

We started with a simple question about a system's memory: supposing I started out in state $|\text{0xC0FFEE}\rangle$, does the system remain recognizably in that state? If the system is a random paramagnet, it obviously does: $[H, \sigma_j^z] = 0$.

If the system is governed by a random-field Heisenberg model, the situation is more complicated. Then $[H_{\text{RFH}}, \sigma_j^z] \neq 0$, but as we have just seen there are operators τ_j^z not so different from σ_j^z that are conserved.

Instead of worrying about the instantaneous decay $\frac{d}{dt}\sigma_j^z = -i[H_{\text{RFH}}, \sigma_j^z]$, attack the long-time limit directly. Work in the fermion picture, with $n_j = \frac{1}{2}(\sigma_j^z + 1)$. If E_α are the eigenenergies of H_{RFH} , then

$$\lim_{t \rightarrow \infty} n_j(t) = \lim_{t \rightarrow \infty} \sum_{\alpha\beta} e^{-i(E_\alpha - E_\beta)t} |\alpha\rangle \langle \alpha| n_j |\beta\rangle \langle \beta| . \quad (1.26)$$

This is strongly oscillatory, so the outcome could depend strongly on when exactly we take the measurement. Clearer would be to start a measurement after some time t_0 and let the measurement go on for some very long time T :

$$\begin{aligned} \frac{1}{T} \lim_{t_0, T \rightarrow \infty} \int_{t_0}^{t_0+T} n_j(t) &= \sum_{\alpha\beta} |\alpha\rangle \langle \alpha| n_j |\beta\rangle \langle \beta| \frac{1}{T} \lim_{t_0, T \rightarrow \infty} \int_{t_0}^{t_0+T} e^{-i(E_\alpha - E_\beta)t} \\ &= \sum_{\alpha} \text{tr} [|\alpha\rangle \langle \alpha| n_j] . \end{aligned} \quad (1.27)$$

The eigenstates $|\alpha\rangle$ are specified by l-bit occupations, e.g. $|\alpha\rangle = |10011101\dots\rangle$, so eigenstate projections are

$$|\alpha\rangle \langle \alpha| = \prod_{j \in \alpha} \tilde{n}_j \prod_{k \notin \alpha} \tilde{n}_k , \quad (1.28)$$

where $j \in \alpha$ indicates that l-bit j is occupied in state $|\alpha\rangle$, and I write

$$\tilde{n}_j = \frac{1}{2}(\tau_j^z + 1) \sim \frac{1}{\xi} n_j + \sum_{jk} e^{-\xi|j-k|} c_j^\dagger c_k + \dots \quad (1.29)$$

for the occupation operator of the l-bit centered at j . That occupation operator looks like

$$\tilde{n}_j \sim \frac{1}{\xi} n_j^z + \sum_{kl} R_{kl}^j e^{-(|k-j|+|j-l|)/\xi} c_k^\dagger c_l . \quad (1.30)$$

Then

$$\frac{1}{T} \lim_{t_0, T \rightarrow \infty} \int_{t_0}^{t_0+T} n_j(t) = \sum_{\alpha} \text{tr} \left[\prod_{k \in \alpha} \tilde{n}_k \prod_{k' \notin \alpha} \tilde{n}_{k'} n_j \right] |\alpha\rangle \langle \alpha| . \quad (1.31)$$

Since

$$n_j \sim \frac{1}{\xi} \tilde{n}_j + [\text{multiple l-bit flips with small coeff's}] , \quad (1.32)$$

this is roughly

$$\begin{aligned} \frac{1}{T} \lim_{t_0, T \rightarrow \infty} \int_{t_0}^{t_0+T} n_j(t) &= \frac{1}{\xi} \sum_{\alpha} \text{tr} \left[\prod_{k \in \alpha} \tilde{n}_k \prod_{k' \notin \alpha} \tilde{n}_{k'} n_j \right] |\alpha\rangle\langle\alpha| \\ &\simeq \tilde{n}_j : \end{aligned} \quad (1.33)$$

the long-time expectation value is given by the nearest l-bit.

1.1.3 Level spacing

This long-time memory, while conceptually satisfying, is too blunt a knife either for diagnostics or for detailed considerations of dynamics. Level spacing statistics—in particular, the distribution of gaps between nearest-in-energy eigenstates, is a much finer tool: it is a standard diagnostic of thermalization and localization, and we will use it extensively in Part II.

To understand the level-spacing statistics of a disordered Hamiltonian, think once again about a random-field Heisenberg model, this time on a system of size L with closed boundaries. To estimate the distribution of smallest gaps, start in an eigenbasis of the underlying Anderson-localized system. Once again, we'll proceed by perturbation theory, but along a different route.

Suppose for the moment that the system is many-body localized, and take two eigenstates $|\alpha\rangle, |\alpha+1\rangle$ of $H_{\&}$ with energies $E_{\&\alpha}, E_{\&\alpha+1}$. Nearest-in-energy eigenstates will differ on most of the system's Anderson l-bits, so on the one hand

$$\delta_{\&} \equiv E_{\&\alpha+1} - E_{\&\alpha} \sim 2^{-L} , \quad (1.34)$$

and on the other

$$\lambda \equiv \langle\alpha| \sigma_j^z \sigma_{j+1}^z |\alpha+1\rangle \sim 2^{-L/2} e^{-L/\zeta} . \quad (1.35)$$

We can then write an effective two-level Hamiltonian on these two levels

$$H = \begin{bmatrix} \delta_{\&} & \lambda \\ \lambda & -\delta_{\&} \end{bmatrix} . \quad (1.36)$$




Figure 1.2: Heuristic derivation of level-spacing distribution

The eigenvalues of (1.36) are

$$E = \pm h \sqrt{\delta_{\&}^2 + \lambda^2}, \quad (1.37)$$

so the gap is $\delta = \sqrt{\delta_{\&}^2 + \lambda^2}$.

What is probability distribution of gaps δ ? Suppose now that the system is MBL, and indeed far from the transition. Both $\delta_{\&}$ and λ are uniformly distributed on their respective scales, but $\lambda \ll \delta_{\&}$: this is the condition (1.20) for the system to be MBL. To see the distribution of δ , Imagine plotting $\delta_{\&}, \lambda$ in a plane (cf Fig. 1.2). For small δ (i.e. $\delta \lesssim \lambda$) the ensemble-averaged level-spacing distribution is

$$\lim_{s \rightarrow 0} p(\delta) d\delta \sim \delta d\delta, \quad \delta \lesssim \lambda; \quad (1.38)$$

for $\delta \gtrsim \lambda$,

$$p(\delta) d\delta \sim \text{const.} \times d\delta, \quad \lambda \lesssim \delta \lesssim h. \quad (1.39)$$

We are therefore unlikely to see gaps smaller than

$$\lambda \sim 2^{-L/2} e^{-L/\zeta}. \quad (1.40)$$

This logic was for gaps $\delta \ll 2^{-L}$. If the gap is larger, $\delta \sim 2^{-L}$, the reasoning above still holds—but $\delta \sim \delta_{\&}$, and you must concern yourself with the fact that $\delta_{\&}$ is not uniform. Take two Anderson levels separated by a gap $\delta'_{\&}$. As $\delta'_{\&}$ increases, the two levels with which you have concerned yourself are increasingly unlikely to be nearest-neighbors in spectrum: there is likely to be a level in between. Taking this into account, one finds that

$$p(\delta) d\delta \simeq p(\delta_{\&}) d\delta_{\&} = \mu e^{-\mu \delta_{\&}} d\delta_{\&}, \quad (1.41)$$

where $\mu = 2^L$ is [average Anderson gap] $^{-1}$. (In Part II, we will frequently concern ourselves with the detail that the rate μ can vary with energy.) So

$$p(\delta) \sim \begin{cases} \delta & \delta \lesssim \lambda \\ \mu e^{-\delta} & \delta \gtrsim \lambda \end{cases}. \quad (1.42)$$

It is convenient to have a single parameter to characterize this level repulsion. Your first thought is probably to compute moments of the level-spacing distribution (1.42), but none are suitable: $\langle s \rangle = 1$ by construction, and higher moments capture the large- s behavior—not the small- s behavior we’re after. Moreover, estimating the level-spacing distribution requires that finicky normalization by the density of states.

Instead, it is conventional to use the so-called *gap ratio*, frequently called the *r-parameter*. Given two adjacent level spacings $\delta_j = E_j - E_{j-1}$ and $E_{j+1} - E_j$, the gap ratio is

$$r_j = \frac{\min(\delta_j, \delta_{j+1})}{\max(\delta_j, \delta_{j+1})}. \quad (1.43)$$

The gap ratio characterizes the distribution (1.42), so one should average it both over levels of a single Hamiltonian and over Hamiltonians drawn from whatever ensemble you’re studying.

To get some intuition, consider two limiting cases:

- If the system has an exact two-fold degeneracy (e.g. a global spin-flip) then small gaps are frequent, and

$$\begin{aligned} \min(\delta_j, \delta_{j+1}) &= 0 \\ \max(\delta_j, \delta_{j+1}) &\neq 0 \end{aligned} \quad (1.44)$$

for

$$r = 0. \quad (1.45)$$

- If the eigenvalues are uniformly spaced then small gaps *never* show up, and

$$\delta_j = \delta_{j+1} \quad (1.46)$$

so $\min(\delta_j, \delta_{j+1}) = \max(\delta_j, \delta_{j+1})$ for

$$r = 1. \quad (1.47)$$

1.1.4 Local dynamics at finite time

We can use the level spacing to tease out details of the intermediate-time dynamics of MBL systems. Think about the evolution of the Heisenberg operator

$$\sigma^z(t). \tag{1.48}$$

How many sites does the operator touch? That is, how many sites in the initial state does $\langle \sigma_j^z(t) \rangle$ depend on? This is essentially a time-dependent correlation length: $\sigma_j^z(t)$ and $\sigma_k^z(t)$ will be correlated to the extent that they depend on the same sites in the initial state.

To compute this correlation function, note that on a timescale t the dynamics is insensitive to gaps smaller than t^{-1} : if two levels are separated by a smaller gap they will not have had time to dephase. All levels differing on a subsystem of diameter

$$l(t) \sim \frac{\ln t}{\ln 2 - \zeta^{-1}} \sim \xi \ln t, \tag{1.49}$$

then, will have had time to dephase, but many levels separated on longer length scales will not—so $\sigma^z(t)$ will have support $l(t) \sim \xi \ln t$. (Cf [109, 158].) This in turn predicts that the entanglement entropy should be

$$S \sim ls \propto \ln t, \tag{1.50}$$

where s is an entropy density; this is exactly what MPS simulations see [16, 194].

This reasoning will make a reappearance in Chs. 6 and 7.

1.2 Matrix product states

The long-time state of an MBL system is characterized by the expectation values of its conserved quantities, the l-bits. l-bits are local, so expectation values of physical operators *here* are only weakly dependent on values of l-bits *there*, and the long-time state should in some sense be “simple”. Moreover, as we just saw, an MBL system that starts out uncorrelated stays uncorrelated for a long time—its correlation length grows as $\ln t$. There should be some nice numerical representations of the kinds of state that result from time evolution with MBL Hamiltonians.

So take an arbitrary state

$$|\psi\rangle = \sum_{s_1 \dots s_N} \psi_{s_1 \dots s_N} |s_1 \dots s_N\rangle \quad (1.51)$$

with short-range correlations on an N -site system with closed boundary conditions. We want a “computationally efficient” representation for the state—that is, an approximation to $|\psi\rangle$ such that that

1. the representation requires little memory, as long as $|\psi\rangle$ has a short correlation length;
2. typical operations require little time (polynomial in system size), again as long as $|\psi\rangle$ has a short correlation length;
3. the representation can “expand” to represent even states that are *not* short-range entangled.

Such a representation exists: the **matrix product state** or **MPS**. The standard reference on matrix product states is the review of Schollwöck [187]. The field has seen many new developments since that review (some described in this thesis), but none have achieved the widespread acceptance of DMRG or TEBD have. For a very different perspective, see the book of Xie Chen et al. [242].

To get some intuition for matrix product states work, start with a **Schmidt decomposition**. Imagine cutting the state in two at some bond l (that is, the bond to the right side of site l). One slightly roundabout way to find the reduced density matrix of the left part of the system is to take the tensor $\psi_{s_1 \dots s_N}$ of Eq. (1.51), group indices to turn it into a matrix,¹ and do a **singular value decomposition (SVD)**

$$\psi_{(s_1 \dots s_l)(s_{l+1} \dots s_N)} = \sum_{\alpha} U_{(s_1 \dots s_l)\alpha} s_{\alpha} V_{(s_{l+1} \dots s_N)} . \quad (1.52)$$

This gives

$$|\psi\rangle = \sum_{\alpha} |x_{L\alpha}\rangle s_{\alpha} |x_{R\alpha}\rangle \quad (1.53)$$

¹In this case, “grouping indices” is exactly the **reshape** operation in NumPy or Julia. In other cases, “grouping indices” will correspond to **reshape** preceded by a reshuffling of indices: `np.transpose` or Julia’s `permutedims`.

with

$$\begin{aligned} |x_{L\alpha}\rangle &= \sum_{s_1 \dots s_l} U_{(s_1 \dots s_l)\alpha} |s_1 \dots s_l\rangle \\ |x_{R\alpha}\rangle &= \sum_{s_{l+1} \dots s_N} V_{(s_{l+1} \dots s_N)\alpha} |s_{l+1} \dots s_N\rangle \end{aligned} \quad (1.54)$$

Since U, V are **isometries**

$$\langle x_{L\alpha} | x_{R\beta} \rangle = \langle x_{L\alpha} | x_{R\beta} \rangle = \delta_{\alpha\beta} . \quad (1.55)$$

(Eq. (1.53) is therefore a Schmidt decomposition.) The reduced density matrices left and right of bond l are then immediately

$$\begin{aligned} \rho_L &= \sum_{\alpha} s_{\alpha}^2 |x_{L\alpha}\rangle \langle x_{L\alpha}| \\ \rho_R &= \sum_{\alpha} s_{\alpha}^2 |x_{R\alpha}\rangle \langle x_{R\alpha}| . \end{aligned} \quad (1.56)$$

So the singular values s_{α} are (the square roots of the) eigenvalues of the density matrices $\rho_{\{L,R\}}$, and entirely specify the state's entanglement properties: the 2nd Rényi and von Neumann entanglement entropies, for example, are

$$\begin{aligned} S_2 &= \sum_{\alpha} s_{\alpha}^4 \\ S_N &= \sum_{\alpha} 2s_{\alpha}^2 \lg s_{\alpha} . \end{aligned} \quad (1.57)$$

If the state has low entanglement entropy at bond l , it should require only a few Schmidt vectors to capture—that is, we should be able to write a pretty good approximation by truncating the sum in Eq. (1.53)

If we follow that line of attack, we're still faced with the problem of representing the states $|x_{\{L,R\}\alpha}\rangle$. What we require is a generalization of Eq. (1.53) that in some sense offers Schmidt decompositions at every bond simultaneously.

A matrix product state offers exactly this property. It consists of a decomposition of the tensor $\psi_{s_1 \dots s_N}$ into a product of matrices (hence the name)

$$\psi_{s_1 \dots s_N} = A_1^{s_1} A_2^{s_2} \dots A_N^{s_N} \quad (1.58)$$

where each of the $A_j^{s_j}$ is a matrix of dimension $\chi \times \chi$ for some χ not too large (except for $A_1^{s_1}$ and $A_N^{s_N}$, which are $1 \times \chi$ and $\chi \times 1$ respectively). χ is called

the **bond dimension** of the MPS; it controls the computational properties and the entanglement of the MPS. The state is

$$|\psi\rangle = \sum_{s_1 \dots s_N} A_1^{s_1} A_2^{s_2} \dots A_N^{s_N} |s_1 \dots s_N\rangle \quad (1.59)$$

(relabeling $\alpha \mapsto \alpha_1$ —in the future such relabelings will be silent) for

$$\psi_{s_1(s_2 \dots s_N)} = A_\alpha^{s_1} r_\alpha V_{\alpha(s_2 \dots s_N)} . \quad (1.60)$$

The index α_1 has dimension χ_1 given by the rank of the $d \times d^{N-1}$ matrix $\psi_{s_1(s_2 \dots s_N)}$. Now apply roughly the same process to rV : regroup indices so that it is a $(\chi_1 d) \times d^{N-2}$ and perform an SVD

$$r_\alpha V_{(\alpha_1 s_2)(s_3 \dots s_N)} = U_{(\alpha_1 s_2) \alpha_2} r_{\alpha_2} V_{\alpha_2(s_3 \dots s_N)} . \quad (1.61)$$

Identify

$$A_{\alpha_1 \alpha_2}^{s_2} = U_{(\alpha_1 s_2) \alpha_2} \quad (1.62)$$

for

$$\psi_{s_1 s_2 \dots s_N} = \sum_{\alpha_1 \alpha_2} A_{\alpha_1}^{s_1} A_{\alpha_1 \alpha_2}^{s_2} r_{\alpha_2} V_{\alpha_2(s_3 \dots s_N)} . \quad (1.63)$$

The index α_2 has dimension χ_2 given by the rank of the $\chi_1 d \times d^{N-1}$ matrix $r_\alpha V_{(\alpha_1 s_2)(s_3 \dots s_N)}$. Proceed in this way until you find (after doing $N - 1$ SVDs)

$$\psi_{s_1 s_2 \dots s_N} = \sum_{\alpha_1 \alpha_2 \dots \alpha_{N-1}} A_{\alpha_1}^{s_1} A_{\alpha_1 \alpha_2}^{s_2} \dots A_{\alpha_{N-2} \alpha_{N-1}}^{s_{N-1}} V_{\alpha_{N-1} s_N} ; \quad (1.64)$$

call $V_{\alpha_{N-1} s_N} \equiv A_{\alpha_{N-1}}^{s_N}$, so

$$\psi_{s_1 s_2 \dots s_N} = \sum_{\alpha_1 \alpha_2 \dots \alpha_{N-1}} A_{\alpha_1}^{s_1} A_{\alpha_1 \alpha_2}^{s_2} \dots A_{\alpha_{N-2} \alpha_{N-1}}^{s_{N-1}} A_{\alpha_{N-1}}^{s_N} \quad (1.65)$$

—this is exactly the form required for Eq. (1.59).

We can bound the bond dimension inductively. χ_1 is the rank of a $d \times d^{N-1}$ matrix, so

$$\chi_1 \leq d, \quad \chi_1 \leq d^{N-1} . \quad (1.66)$$

χ_2 is the rank of a $(\chi_1 d) \times d^{N-2}$ matrix, so

$$\chi_2 \leq (\chi_1 d) \leq d^2, \quad \chi_2 \leq d^{N-2} . \quad (1.67)$$

Proceeding in this way

$$\chi_n \leq d^n, \quad \chi_n \leq d^{N-n}, \quad (1.68)$$

and in general

$$\chi_n \leq d^{\lfloor n/2 \rfloor}. \quad (1.69)$$

In fact one rarely constructs an MPS in this way. The utility of MPSs is in the fact that they are efficient representations of states too large to represent otherwise. If the system size is small enough that one can store the whole coefficient tensor $\psi_{s_1 \dots s_N}$ you are almost certainly better off working with the exact state in an ordinary representation. But the fact that all states can be represented exactly is useful. By taking large bond dimensions, one can perform exact simulations and use them to check your low-bond dimension approximate simulations (we will do this over and over in Part I). More generally, this motivates convergence testing. Suppose you are using an approximation method that cuts MPSs down to some low bond dimension cap. If doubling (say) that bond dimension cap doesn't change the results, you can argue that—since the method will give the exact result for sufficiently large bond dimension cap—this insensitivity means that you have the right answer.

1.2.1 Graphical notation for tensors and tensor contractions

Writing out indices in sums like (1.65) is flexible but cumbersome, while the matrix-multiplication form of Eq. (1.59) is compact but inflexible. It is customary to use a compact graphical notation for tensors and tensor contractions. Tensors are represented by boxes with legs, each leg representing a tensor index, so for example

$$A_{\alpha_1 \alpha_2}^{s_2} = \begin{array}{c} \square \\ \text{---} \text{---} \\ | \\ \text{---} \end{array}. \quad (1.70)$$

The mapping between leg and index is generally to be understood from context and implicit convention; in this case, the left and right legs represent α_1, α_2 respectively, and the bottom leg s_2 . Frequently the box is omitted, and a tensor is represented by a vertex:

$$A_{\alpha_1 \alpha_2}^{s_2} = \begin{array}{c} \square \\ \text{---} \text{---} \\ | \\ \text{---} \end{array} = \begin{array}{c} \text{---} \text{---} \\ | \\ \text{---} \end{array}. \quad (1.71)$$

Tensor contraction—in which one sets indices equal and sums—is indicated by joining the lines corresponding to the indices, so

$$\sum_{\alpha_1} A_{\alpha_1}^{s_1} A_{\alpha_1 \alpha_2}^{s_2} = \begin{array}{c} \text{---} \\ \text{---} \\ \text{---} \end{array} . \quad (1.72)$$

In this notation the coefficient tensor ψ is

$$\psi_{s_1 s_2 \dots s_N} = \sum_{\alpha_1 \alpha_2 \dots \alpha_{N-1}} A_{\alpha_1}^{s_1} A_{\alpha_1 \alpha_2}^{s_2} \dots A_{\alpha_{N-2} \alpha_{N-1}}^{s_{N-1}} A_{\alpha_{N-1}}^{s_N} = \begin{array}{c} \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \end{array} . \quad (1.73)$$

Complex conjugation is represented by flipping the diagram top-to-bottom

$$\psi_{s_1 s_2 \dots s_N}^* \begin{array}{c} \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \end{array} , \quad (1.74)$$

so a state's normalization is

$$\langle \psi | \psi \rangle = \sum_{s_1 \dots s_N} \psi_{s_1 s_2 \dots s_N} \psi_{s_1 s_2 \dots s_N}^* = \begin{array}{c} \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \end{array} . \quad (1.75)$$

1.2.2 Expectation values and correlations in MPSs

One is not typically satisfied with computing the normalization of a state. Suppose one has an onsite operator B with matrix elements $B_{ss'}$. Then

$$B_{ss'} = \begin{array}{c} \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \end{array} . \quad (1.76)$$

and

$$\langle \psi | B | \psi \rangle = \begin{array}{c} \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \end{array} . \quad (1.77)$$

(Ordinarily the operators are Hermitian; if they are not, some care must be taken in interpreting diagrams like this one.)

This is less difficult to evaluate than it may appear. One can perform the contractions in a diagram like Eq. (1.77) by “sweeping in” from left and right—this is much more efficient than forming $\psi_{s_1 \dots s_N}$ explicitly. Reordering the contractions corresponds to reshuffling sums and multiplications in an explicit representation analogous to Eq. (1.65), so one can iteratively construct left

and right environment tensors

$$\begin{aligned}
 E_{Lj} &\equiv \left[\begin{array}{c} \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \end{array} \right] = \left[\begin{array}{c} \square \\ \square \\ \square \\ \square \\ \square \end{array} \right] \\
 E_{Lj} &= \left[\begin{array}{c} \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \end{array} \right] \leftarrow \left[\begin{array}{c} \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \end{array} \right] \\
 E_{Rj} &\equiv \left[\begin{array}{c} \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \end{array} \right] = \left[\begin{array}{c} \square \\ \square \\ \square \\ \square \\ \square \end{array} \right] \\
 E_{Lj} &= \left[\begin{array}{c} \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \end{array} \right] \leftarrow \left[\begin{array}{c} \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \end{array} \right]
 \end{aligned} \tag{1.78}$$

Then the expectation value is

$$\langle \psi | B | \psi \rangle = \left[\begin{array}{c} \square \\ \square \\ \square \\ \square \\ \square \end{array} \right] \left[\begin{array}{c} \square \\ \square \\ \square \\ \square \\ \square \end{array} \right] \tag{1.79}$$

Multi-site operators are an easy generalization: if O is a two-site operator, it has matrix elements

$$O_{ss'} = \left[\begin{array}{c} \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \end{array} \right] \tag{1.80}$$

and expectation value

$$\langle \psi | B | \psi \rangle = \left[\begin{array}{c} \square \\ \square \\ \square \\ \square \\ \square \end{array} \right] \left[\begin{array}{c} \square \\ \square \\ \square \\ \square \\ \square \end{array} \right] \tag{1.81}$$

None of these computations require constructing the coefficient tensor ψ , which would require memory (and time) exponential in system size. Rather, they require time and memory linear in system size and polynomial in the maximum bond dimension χ , and if you want many expectation values you can store the environment tensors and amortize the cost of computing them.

1.2.3 Canonical forms and orthogonality centers

The time-consuming part of the expectation value constructions of Sec. 1.2.2 is the computation of the environment tensors of Eq. (1.78). But if we choose the A matrices to have appropriate properties, then computing the environment tensors (1.78) becomes immediate.

Suppose that the MPS is constructed as in Sec. ???. Then, because the A matrices are (for $j < N$) the left unitary matrices of SVDs, we have

$$\sum_{s_j \alpha_{j-1}} A_{\alpha_{j-1} \alpha_j}^{s_j} A_{\alpha_{j-1} \alpha'_j}^{s_j} = \delta_{\alpha_j \alpha'_j}, \quad j < N \tag{1.82}$$

or in graphical notation

$$\left[\begin{array}{c} \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \end{array} \right] = \left[\begin{array}{c} \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \end{array} \right], \tag{1.83}$$

where the curved line is $\zeta = \delta_{\alpha \alpha'}$.

So

$$\begin{aligned}
E_{L1} &= \boxed{} = \text{C} \\
E_{L2} &= \boxed{} = \text{C} \\
E_{Ln} &= \boxed{} = \text{C}.
\end{aligned} \tag{1.84}$$

To find the analogous result for the right side, return to decompositions of the form Eq. (1.63). After $n - 1$ steps we have

$$\psi_{s_1 s_2 \dots s_N} = \sum_{\alpha_1 \alpha_2} A_{\alpha_1}^{s_1} A_{\alpha_1 \alpha_2}^{s_2} \dots A_{\alpha_{n-1} \alpha_n}^{s_{n-1}} r_{\alpha_n} V_{\alpha_n(s_n \dots s_N)} \tag{1.85}$$

so, comparing with Eq. (1.65),

$$\sum_{\alpha_{n+1} \dots \alpha_{N-1}} A_{\alpha_n \alpha_{n+1}}^{s_n} \dots A_{\alpha_{N-1}}^{s_N} = \sum_{\alpha_n} r_{\alpha_n} V_{\alpha_n(s_n \dots s_N)} \tag{1.86}$$

and

$$E = \boxed{} = \sum_{s_n \dots s_N} r_{\alpha_n} r_{\alpha'_n} V_{\alpha_n(s_n \dots s_N)}^* V_{\alpha'_n(s_n \dots s_N)} = r_{\alpha}^2 \delta_{\alpha \alpha'}. \tag{1.87}$$

Working from right to left one can deduce the right-side analogue of the local condition Eq. (1.83):

$$\sum_{s_j \alpha_j} A_{\alpha_{j-1} \alpha_j}^{s_j^*} A_{\alpha'_{j-1} \alpha_j}^{s_j} = r_{\alpha_{j-1}}^2 \delta_{\alpha_{j-1} \alpha'_{j-1}}, \quad j < N \tag{1.88}$$

or in graphical notation

$$\mathbb{D} = \text{◇}. \tag{1.89}$$

The diamond conventionally indicates a diagonal matrix, here with diagonal entries $r_{\alpha_j}^2$. Now a single-site expectation value is

$$\boxed{} = \boxed{\text{◇}}. \tag{1.90}$$

MPSs with the property Eq. (1.83) are called **left canonical**. If a state is called “left canonical” one generally expects it also to have the property (1.90), but arranging this can require some effort. One can define **right canonical**

MPSs analogously: these are MPSs

$$|\psi\rangle = \sum_{s_1 \dots s_N} B_1^{s_1} B_2^{s_2} \dots B_N^{s_N} |s_1 \dots s_N\rangle . \quad (1.91)$$

such that

$$\sum_{s_j \alpha_{j-1}} B_{\alpha_{j-1} \alpha_j}^{s_j^*} B_{\alpha_{j-1} \alpha'_j}^{s_j} = r_{\alpha_j \alpha'_j}^2 , \quad (1.92)$$

and typically

$$\sum_{s_j \alpha_j} B_{\alpha_{j-1} \alpha_j}^{s_j^*} B_{\alpha_{j-1} \alpha'_j}^{s_j} = \delta_{\alpha_{j-1} \alpha'_{j-1}} , \quad (1.93)$$

$j < N$, or

$$\mathbb{D} = \mathbb{J} \quad (1.94)$$

Expectation values are then

$$\langle \psi | |B | \psi \rangle = \langle \mathbb{D} \rangle . \quad (1.95)$$

How does one switch between left and right canonical forms? One's first thought is to write

$$B_{\alpha_{j-1} \alpha_j}^{s_j} = (r_{\alpha_{j-1}}^{(j-1)})^{-1} A_{\alpha_j \alpha_{j-1}}^{s_j} r_j^{(j)} , \quad (1.96)$$

(where I write $r_{\alpha}^{(j-1),(j)}$ to emphasize that these are two different sets of singular values r). This gives the same state, because the $r^{(j-1)}$ on the right side of $B^{(j-1)}$ will multiply the $[r^{(j-1)}]^{-1}$ on the left side of $B^{(j)}$. In principle this gives the desired result, but it is numerically unstable because it involves dividing by the singular values $r_{\alpha_{j-1}}^{(j-1)}$, which generically range in magnitude from ~ 1 to numerical precision.

Better is to sweep from right to left, performing SVDs at each site. Start with a left-canonical MPS

$$|\psi\rangle = \sum_{s_1 \dots s_N} A_1^{s_1} A_2^{s_2} \dots A_N^{s_N} |s_1 \dots s_N\rangle . \quad (1.97)$$

SVD

$$A_{\alpha_N s_N} = U_{\alpha_N \alpha'_N} r_{\alpha'_N} V_{\alpha'_N s_N} . \quad (1.98)$$

Identify

$$B_{\alpha'_N}^{s_N} \equiv V_{\alpha'_N s_N} ; \quad (1.99)$$

now

$$|\psi\rangle = \sum_{s_1 \dots s_N} A_1^{s_1} A_2^{s_2} \dots A_{N-2}^{s_{N-2}} A_{N-1}^{s_{N-1}} U r B_N^{s_N} |s_1 \dots s_N\rangle , \quad (1.100)$$

where B_N satisfies the right canonical relation (1.92) Now multiply that $U r$ into the A to its left, defining

$$X_{\alpha_{N-2} \alpha_{N-1}}^{s_{N-1}} = A_{\alpha_{N-2} \alpha_{N-1}}^{s_{N-1}} r_{\alpha_{N-1}} . \quad (1.101)$$

SVD

$$X_{\alpha_{N-2}(\alpha_{N-1} s_{N-1})} = U_{\alpha_{N-2} \gamma_{N-2}} r_{\gamma_{N-2}} V_{\gamma_{N-2}(\alpha_{N-1} s_{N-1})} . \quad (1.102)$$

Identify

$$B_{\alpha_{N-2} \alpha_{N-1}}^{s_{N-1}} = V_{\alpha_{N-2}(\alpha_{N-1} s_{N-1})} ; \quad (1.103)$$

this B again satisfies the right canonical relation (1.92).

$$|\psi\rangle = \sum_{s_1 \dots s_N} A_1^{s_1} A_2^{s_2} \dots A_{N-2}^{s_{N-2}} U r B_{N-1}^{s_{N-1}} B_N^{s_N} |s_1 \dots s_N\rangle . \quad (1.104)$$

Proceeding in this way, after $N - n$ SVDs

$$|\psi\rangle = \sum_{s_1 \dots s_N} A_1^{s_1} A_2^{s_2} \dots A_n^{s_n} U r B_{n+1}^{s_{n+1}} \dots B_N^{s_N} |s_1 \dots s_N\rangle \quad (1.105)$$

and after $N - 1$

$$|\psi\rangle = \sum_{s_1 \dots s_N} B_1^{s_1} \dots B_N^{s_N} |s_1 \dots s_N\rangle \quad (1.106)$$

with all the B satisfying (1.92): we have put the MPS into right canonical form, as desired, without ever dividing by singular values.

A state in the intermediate form (1.105) is said to have a (zero-site) **orthogonality center** at bond n^2 , or to be in **mixed canonical form**. This is equivalent to the Schmidt decomposition.

²One sometimes uses “zero-site” to distinguish MPSs like (1.105) from those in which the $U r$ has been multiplied into the A to its left; this is called a “one-site” orthogonality center. I do not use one-site orthogonality centers in this thesis, so I do not say “zero-site”.

1.3 Dynamics and matrix product states

In Sec. 1.1 we asked what happened to a state

$$|0\text{xCOFFEE}\rangle = |\uparrow\uparrow\downarrow\downarrow\downarrow\downarrow\downarrow\downarrow\uparrow\uparrow\uparrow\uparrow\uparrow\uparrow\uparrow\uparrow\uparrow\downarrow\uparrow\uparrow\uparrow\downarrow\rangle, \quad (1.107)$$

under evolution by a Hamiltonian

$$H_{\text{RFH}} = \sum \vec{\sigma}_j \cdot \vec{\sigma}_{j+1} + \sum_j h_j \sigma_j^z. \quad (1.108)$$

We found that at long times the system retained some memory of that initial state $|0\text{xCOFFEE}\rangle$, and that it develops correlations slowly—with a correlation length $l \sim \xi \log t$, ξ the localization length of H_{RFH} . Then I argued (briefly) that matrix product states are natural tools for numerically simulating MBL, and described (at length) natural operations on those matrix product states. Now I need to describe one more “natural operation”—time evolution.

A Hamiltonian like H_{RFH} acts locally. If I time evolve $|0\text{xCOFFEE}\rangle$ by H_{RFH} for some time $\delta t \ll 1$, then it is unreasonable to expect that the state of site 1 will come to depend on the state of site 20 in any substantial way—except to the extent that they depended on each other to start out with. Hamiltonian does not quickly introduce new correlations. This is doubly true of an MBL system.

We can attempt to quantify the notion that the Hamiltonian does not quickly introduce long-distance correlations by noticing that

$$e^{-iH\delta t} = 1 - iH\delta t + O(H\delta t^2) \quad (1.109)$$

is local. Eq. 1.109 is of dubious suitability for describing dynamics, though, because not unitary. We can remedy this defect, while preserving insight into locality, by splitting the Hamiltonian up into even and odd terms and time evolving first by the odd terms, and then by the even terms. Write two-site gates

$$F_j(\delta t) = \exp \left[-i\delta t \left(\frac{1}{2}(h_j^z + h_{j+1}^z) + \vec{\sigma}_j \cdot \vec{\sigma}_{j+1} \right) \right]. \quad (1.110)$$

Then—as one can see by comparing Taylor series—

$$e^{-iH_{\text{RFH}}\delta t} = U_{\delta t} + O(\delta t^2) \quad U_{\delta t} \equiv \prod_{j=1}^{L/2} [F_{2j}(\delta t)] \prod_{j=1}^{L/2} [F_{2j+1}(\delta t)] \quad (1.111)$$

and

$$e^{-iH_{\text{RFH}}t} = U_{\delta t}(t) + O(\delta t), \quad U_{\delta t}(t) \equiv [U_{\delta t}]^{t/\delta t} + O(\delta t). \quad (1.112)$$

There exist Trotter decomposition's with faster convergence, e.g. the second-order decomposition

$$e^{-iH_{\text{RFH}}\delta t} = U_{\delta t}^{(2)} + O(\delta t^3), \quad (1.113)$$

$$U_{\delta t}^{(2)} = \prod_{j=1}^{L/2} [F_{2j}(\delta t/2)] \prod_{j=1}^{L/2} [F_{2j+1}(\delta t)] \prod_{j=1}^{L/2} [F_{2j}(\delta t/2)]$$

which takes little more time than (1.112) to compute, thanks to a convenient telescoping property.

This higher-order convergence is of secondary (practical) importance for times $t \gtrsim 1$ in “natural units” (some coupling). In fact the error estimate $O(\delta t)$ in Eq. (1.112) is extraordinarily pessimistic: one can show that $U_{\delta t}(t)$ approximates $e^{-iH_{\text{RFH}}t}$ well to times exponentially large in $1/\delta t$ [4, 6, 7, 57, 126, 155].

Alternatively, we can decompose

$$e^{-iH\delta t} \simeq \underbrace{\prod_{j=1}^{L-1} [F_j(\delta t/2)]}_{\text{leftward sweep after}} \underbrace{\prod_{j=L-1}^1 [F_j(\delta t/2)]}_{\text{rightward sweep}}; \quad (1.114)$$

the error is again of order $L\|h_j\|^3\delta t^3$, where $\|h_j\|$ is an estimate of the typical magnitude of the terms h_j . We call this a “boustrophedon” Trotter decomposition. If you wish to time evolve an MPS in mixed-canonical form and truncate at the orthogonality center, as we will in 2, this makes it easy to move the orthogonality center.

1.4 Matrix product operators

In Sec. 1.3, I represented time evolution in terms of Trotter gates—that is, of local unitaries. This is a restrictive representation: if you wish to work with more general operators than local Hamiltonians, like long-range Hamiltonians or density operators, you need a matrix product representation of an operator.

To write a **matrix product operator (MPO)** representation of some generic

operator O , decompose it in a basis of local operators

$$O = \sum_{\mu_1 \dots \mu_N} O_{\mu_1 \dots \mu_N} \sigma^{\mu_1} \dots \sigma^{\mu_N} , \quad (1.115)$$

and decompose the tensor $O_{\mu_1 \dots \mu_N}$ as a product of matrices

$$O_{\mu_1 \dots \mu_N} = A^{\mu_1} \dots A^{\mu_N} . \quad (1.116)$$

All of the tricks for manipulating MPSs (gauge transformations, canonical forms, truncations, etc.) apply to MPOs, but frequently with altered interpretation. The operators are vectors in a larger Hilbert space than the system's physical Hilbert space—but their inner product is

$$\langle A|B \rangle = \text{tr } A^\dagger B . \quad (1.117)$$

1.4.1 MPO representation of an exponential interaction

Take a (portion of a) Hamiltonian

$$H_{\text{exp}} = \sum_{i < j} \sigma_i^z \sigma_j^z e^{-s|j-i|} \quad (1.118)$$

for some constant s . We can build this Hamiltonian by walking in from the right end of the chain with the recursion relation

$$\begin{bmatrix} H_{R_i} \\ h_{R_i} \\ I_{R_i} \end{bmatrix} = \begin{bmatrix} I_i & \sigma_i^x & 0 \\ 0 & e^{-s} I_i & \sigma_i^x \\ 0 & 0 & I_i \end{bmatrix} \otimes \begin{bmatrix} H_{R_{i+1}} \\ h_{R_{i+1}} \\ I_{R_{i+1}} \end{bmatrix} \quad (1.119)$$

starting with a right-end base case

$$\begin{bmatrix} H_{R_{L+1}} \\ h_{R_{L+1}} \\ I_{R_{L+1}} \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} . \quad (1.120)$$

The Hamiltonian is then H_{R_0} —that is, the thing in the first slot when you've walked all the way to the left end of the chain.

The matrices constituting the MPO are then

$$\begin{aligned}
W_1 &= \begin{bmatrix} I_1 \\ \sigma_1^z \\ 0 \end{bmatrix} \\
W_{1 < i < L} &= \begin{bmatrix} I_i & \sigma_i^z & 0 \\ 0 & s^{-1}I_i & \sigma_i^z \\ 0 & 0 & I_i \end{bmatrix} \\
W_L &= \begin{bmatrix} 0 \\ \sigma_i^z \\ I_i \end{bmatrix}
\end{aligned} \tag{1.121}$$

1.4.2 MPO representation of a power-law Hamiltonian

This is from [241]; cf [240].

Take for concreteness

$$\begin{aligned}
H &= H_{\text{local}} + H_{\text{power}}, \\
H_{\text{power}} &= \sum_{i < j} \frac{\sigma_i^z \sigma_j^z}{|j - i|^q},
\end{aligned} \tag{1.122}$$

where H_{local} is 2-local.

1.4.2.1 Exact representation

To understand how to represent this as an MPO, first note that there exist unique D_k , $k = 1 \dots L$, such that

$$\frac{1}{j^q} = \sum_{k=1}^L C_s e^{-js_k/L}, \quad j \in 1 \dots L, \quad s_k = k/L \tag{1.123}$$

—this simply by linear independence of the vectors $e^{-js_k/L}$.

Then

$$H_{\text{power}} = \sum_k C_k \sum_{i < j} e^{-|j-i|s_k/L} \sigma_i^z \sigma_j^z, \tag{1.124}$$

which has a recursion relation representation

$$\begin{bmatrix} H_{R_i} \\ h_{R_i} \\ I_{R_i} \end{bmatrix} = \begin{bmatrix} I_i & C_1 \sigma_i^z & C_2 \sigma_i^z & \dots & 0 \\ 0 & e^{-s_1} I_i & 0 & \dots & \sigma_i^z \\ \vdots & & \ddots & & \vdots \\ & & & e^{-s_L} I_i & \sigma_i^z \\ 0 & \dots & & 0 & I_i \end{bmatrix} \otimes \begin{bmatrix} H_{R_{i+1}} \\ h_{R_{i+1}} \\ I_{R_{i+1}} \end{bmatrix} \quad (1.125)$$

and resulting MPO representation. The Hamiltonian has an MPO representation given by adding the MPO representations for H_{power} and H_{local} .

1.4.2.2 Approximate representation

The key insight in [241] is that you can get away with many fewer than L exponentials if you're willing to accept small errors in your interaction. This is apparently well-studied in the control system literature under the name model reduction.

The approximation is given by the recursion relation

$$\begin{bmatrix} H_{R_i} \\ h_{R_i} \\ I_{R_i} \end{bmatrix} = W_i \otimes \begin{bmatrix} H_{R_{i+1}} \\ h_{R_{i+1}} \\ I_{R_{i+1}} \end{bmatrix} \quad (1.126)$$

with

$$W_i \equiv \begin{bmatrix} I_i & C_1 \sigma_i^z & C_2 \sigma_i^z & \dots & 0 \\ 0 & e^{-s_1} I_i & 0 & \dots & \sigma_i^z \\ \vdots & & \ddots & & \vdots \\ & & & e^{-s_n} I_i & \sigma_i^z \\ 0 & \dots & & 0 & I_i \end{bmatrix} \quad (1.127)$$

(with resulting MPO representation) where the $n < L$ decay wavevectors $s_1 \dots s_n$ are chosen by the algorithm. Note that the s_k can be complex.



1.4.3 MPDOs

1.4.3.1 Closed-system dynamics of an MPDO

Consider a Hamiltonian

$$H = \sum H_j \quad (1.128)$$

(for instance (2.40)) where H_j is supported on sites $j, j+1$. Call the onsite Hilbert space \mathcal{H}_j and its dual \mathcal{H}_j^* . Pure states then live in a Hilbert space

Figure 1.3: Decay rates s_k for a power law $1/j^{1.5}$ Figure 1.4: Decay lengths s_k^{-1} for a power law $1/j^{1.5}$

$\mathcal{H} = \mathcal{H}_j^{\otimes L}$, and density operators ρ in a Hilbert space $\mathcal{H} \otimes \mathcal{H}^* = [\mathcal{H}_j \otimes \mathcal{H}_j^*]^{\otimes L}$. Closed-system Hamiltonian evolution, then, is

$$\frac{d}{dt}\rho = -i[H, \rho] \equiv -iH^\sharp \rho \quad (1.129)$$

with a linear superoperator Hamiltonian defined

$$\begin{aligned} H^\sharp &: \mathcal{H} \otimes \mathcal{H}^* \rightarrow \mathcal{H} \otimes \mathcal{H}^*, \\ H^\sharp &= H \otimes 1 - 1 \otimes H. \end{aligned} \quad (1.130)$$

For a spin-1/2 chain, we can write $S_j^{x,y,z}$ for operators on the ordinary (“ket”) space \mathcal{H}_j and $T_j^{x,y,z}$ for operators on the dual (“bra”) space \mathcal{H}_j^* ; in this notation, the superoperator corresponding to our fruit-fly Ising Hamiltonian (2.40) is

$$\begin{aligned} H^\sharp &= \sum H_j^\sharp \\ &\equiv \sum \left[(S_j^z S_{j+1}^z - T_j^z T_{j+1}^z) \right. \\ &\quad \left. + \frac{1}{2} h^x (S_j^x - T_j^x) + \frac{1}{2} h^z (S_j^z - T_j^z) \right]. \end{aligned} \quad (1.131)$$

With H^\sharp in hand, we can proceed by Trotter decomposition, as with an MPS.

1.4.4 Matrix product density operator representations of Gibbs states

1.4.4.1 Construction

In sections 2.3.3 and 2.3.4 we will require a Gibbs MPDO as our initial state. The Gibbs state is

$$\rho \propto e^{-\beta H} = e^{-\beta H/2} I e^{-\beta H/2}, \quad (1.132)$$

which is precisely the imaginary-time evolution of the product MPDO I by the Hamiltonian superoperator

$$H_{\text{therm}}^\sharp = H \otimes 1 + 1 \otimes H \quad (1.133)$$

out to time $\beta/2$. We approximate this imaginary-time evolution not by tDMRG with the boustrophedon Trotter decomposition (1.114), but by ordinary TEBD using the trick of Hastings for numerical stability (as described in 7.3.2 of Schollwöck’s magisterial review [187], q.v.).

1.4.4.2 Estimating thermal expectation values

In analyzing the time evolution of ETH states, one naturally requires Gibbs state expectation values as a function of total energy (or, equivalently, energy density): the long time limit of an expectation value is given by its expectation value in a Gibbs state whose energy density matches that of the initial state. We tabulate energy densities and observables of interests for Gibbs states at a variety of temperatures using MPDOs as described above; to find the long-time limit of an expectation value, we measure the energy density of the initial state and linearly interpolate between the two nearest Gibbs energy densities. Note that this does not account for Trotter heating (that is, the fact that—because the system actually simulated is a Floquet system with period given by the Trotter step δt , its energy as measured by the Hamiltonian simpliciter gradually increases).

1.5 Some essentials of quantum thermodynamics

The classical Otto engine consists of a gas that expands, cools, contracts, and heats [179]. During the two isentropic (constant-entropy) strokes, the gas’s volume is tuned between values V_1 and $V_2 < V_1$. The *compression ratio* is defined as $r := \frac{V_1}{V_2}$. The heating and cooling are isochoric (constant-volume).

The engine outputs a net amount W_{tot} of work per cycle, absorbing heat $Q_{\text{in}} > 0$ during the heating isochore.

A general engine's thermodynamic efficiency is

$$\eta := \frac{W_{\text{tot}}}{Q_{\text{in}}}. \quad (1.134)$$

The Otto engine operates at the efficiency

$$\eta_{\text{Otto}} = 1 - \frac{1}{r^{\gamma-1}} < \eta_{\text{Carnot}}. \quad (1.135)$$

A ratio of the gas's constant-pressure and constant-volume specific heats is denoted by $\gamma := \frac{C_P}{C_V}$. The Carnot efficiency η_{Carnot} upper-bounds the efficiency of every thermodynamic engine that involves just two heat baths.

A quantum Otto cycle for harmonic oscillators was discussed in Refs. [3, 51, 53, 102, 123, 190, 215, 244]. The quantum harmonic oscillator's (QHO's) gap plays the role of the classical Otto engine's volume. Let ω and $\Omega > \omega$ denote the values between which the angular frequency is tuned. The ideal QHO Otto cycle operates at the efficiency

$$\eta_{\text{QHO}} = 1 - \frac{\omega}{\Omega}. \quad (1.136)$$

This oscillator model resembles the qubit toy model that informs our MBL Otto cycle (Sec. 6.1).

The heat and work exchanged by slowly tuned systems are defined as

$$W := \int_0^\tau dt \text{Tr} \left(\rho \frac{dH}{dt} \right) \quad \text{and} \quad (1.137)$$

$$Q := \int_0^\tau dt \text{Tr} \left(\frac{d\rho}{dt} H \right) \quad (1.138)$$

in quantum thermodynamics [215]. This Q definition is narrower than the definition prevalent in the MBL literature [39, 44, 70, 137]: here, all energy exchanged during unitary evolution counts as work.

Part I

Thermalizing systems and matrix product density operators

SIMULATING QUANTUM DYNAMICS OF THERMALIZING SYSTEMS

Questions about how (and whether) hydrodynamic behavior emerges from microscopic quantum physics arise frequently in condensed matter physics. The exploration of this physics is hampered by the limitations of existing numerical methods (cf Fig. 2.1). Numerically exact methods (like exact diagonalization and Krylov subspace methods) can treat the dynamical properties of small systems at arbitrary times, but require memory and computation time exponential in system size. Matrix product state methods, on the other hand, can treat large systems—but only when the systems have little entanglement entropy. This means that for thermalizing systems, whose entanglement entropy grows linearly with time, MPS methods can only treat short-time behavior (see Figure 2.1).

This chapter describes a numerical method (“density matrix truncation” or “DMT”) based on matrix product representations of density operators. The al-




Figure 2.1: For small systems, exact diagonalization and related methods can treat time evolution of small thermalizing systems to long times. Matrix product state methods, on the other hand, can treat time evolution of large thermalizing systems, but only to short times, and hydrodynamic effective theories can phenomenologically describe the long-time limit, but not intermediate times. Our method, DMT, can treat large ETH systems at all times.

gorithm can accurately simulate not only short-time, low-entanglement behavior and long-time, hydrodynamic behavior, but also the complex intermediate-time behavior from which the hydrodynamics emerges. While using matrix-product representations of mixed states is not new, the core of our algorithm is a new method for truncating matrix product density operators (MPDOs). This truncation exactly preserves the energy density of the system and other local conserved quantities, with the aim of leaving the hydrodynamics unaffected. It also avoids generating unphysical density matrices with negative entropy. At the same time, it is efficient enough that by taking large (but constant in system size) bond dimension one can capture the thermalization process.

We first (in Sec. 2.1) offer some background on matrix product state methods and intuition for why a method using MPDOs should be able to efficiently simulate time evolution governed by Hamiltonians satisfying the eigenstate thermalization hypothesis (ETH) to arbitrary times. We also motivate certain properties of our method. We then describe (in Sec. 2.2) our algorithm for time evolution. This algorithm consists of the time-evolution framework described in Sec. 1.3 (cf [214]), paired with a novel scheme for truncating MPDOs. We then apply our algorithm to time evolution starting from a pure state (Sec. 2.3.2) and find that it qualitatively improves upon existing methods. Applied to a mixed state (Sec. 2.3.3 and Sec. 2.3.4), we find that our algorithm DMT matches or exceeds the state of the art. We conclude with directions we hope will improve on the method.

2.1 Background and intuition

2.1.1 Background: matrix product state methods

Simulating real-time evolution of many-body quantum-mechanical systems is hard: a system of L sites generically requires storage and computation time exponential in L . One line of attack, e.g. time-evolving block decimation (TEBD) [213, 214, 228], proceeds by representing unitary time evolution as a series of small time-steps applied to a matrix-product state (MPS) representation of a pure state. These matrix-product structures offer efficient representations of certain states (broadly speaking, “low-entanglement states”) in the sense that typical operations require polynomial time and memory.

Matrix product state simulations of time-evolving pure states are stymied by the fast increase of entanglement entropy with time, which grows linearly in time for a typical global quench. When one compresses a pure state as a matrix

product state, memory and computation time requirements grow exponentially in the entanglement entropy among the subsystems, so this linear growth in entanglement entropy sets a hard upper limit on the timescales on which matrix product states are useful, though a variety of methods have been used in attempts to circumvent this limit. [14, 58, 64, 78–81, 90, 157, 204, 223, 231]. One case in which entanglement growth does not limit the useful timescale of matrix product state simulations is the dynamics of many-body localized systems, which exhibit a modest logarithmic entanglement growth [16, 192]. On the thermal side of the localization transition, however, entanglement grows quickly, and even the transition itself is expected to show extensive (volume-law) entanglement [76]—consequently, pure-state time evolution under Hamiltonians remains restricted to short times.

Research into mixed-state time evolution and Lindblad dynamics has also progressed. It has been proven that density matrices (and purifications) of Gibbs states of local Hamiltonians have efficient matrix product representations [85, 154, 163, 247]. Two schools of thought have used this insight to develop a series of methods for simulating time evolution. One school employs density matrices [21, 41, 95, 95, 101, 151, 152, 167, 168, 175–178, 212, 221, 222, 247]. They note that the space of operators on a spin chain is the tensor product of onsite operator spaces, just as the space of many-body pure states being a tensor product on onsite Hilbert spaces; the chief difference (in this view) is merely the dimensionality of the onsite space. For example, on a spin-half chain, the space of onsite operators is four-dimensional, while the space of pure states is two dimensional. This school then applies familiar pure state methods, including the creation and truncation of matrix product states and time evolution by TEBD, to density matrices—which are, after all, vectors in a larger space. The resulting truncation algorithms minimize the error according to the Hilbert-Schmidt (Frobenius) norm. In certain situations—in particular, dynamics near thermal equilibrium or a non-equilibrium steady state—this approach works well. In other situations, however—in particular, time evolution starting from a pure state—the density matrices suffer from a catastrophic loss of positivity. (Even checking positivity is NP-hard in the system size [117].)

The second school [17, 18, 23, 58, 60, 103–106, 208] uses purifications instead of density matrices to represent mixed states. They pair each site in the system with an ancilla, a notional site representing the bath. The mixed nature of the system is represented by entanglement between sites and their ancillae,

so the system and ancillae together are in a pure state. Grouping each site and its ancilla into a larger onsite Hilbert space, one can write a matrix product representation for this pure state and apply the usual methods (truncation, TEBD, etc.) This solves the positivity problem: unlike density matrices, where many operators with reasonable matrix product representations are not positive and hence are invalid as density matrices, every representable vector is a valid state. Moreover, since one can act with a unitary on the space of ancillae without changing the physical state, one can try to exploit this freedom to reduce the matrix product state bond dimension of the purification [17, 104]. There is also a hybrid approach which locally unzips a density matrix into a purification, which preserves positivity by construction [225]. These purification methods employ truncations which minimize error according to the inner product $\langle \cdot | \cdot \rangle$ defined on the whole (system with ancillae) state.

Neither the Frobenius norm on density matrices nor the quantum-mechanical norm on purifications is the correct notion of error. In the case of density matrices, the Frobenius norm fails to account for the fact that truncations that change the component of the density-matrix vector along the identity (i.e. which are not trace-preserving) are disastrous, because they can lead to loss of positivity. Moreover, neither notion of error captures spatial locality: a good notion of error should prioritize short-range properties of the model and guarantee that certain quantities (the local conserved quantities of the model under consideration, like energy density or spin) are unchanged. Since the methods of both the density-matrix and purification schools generically change the model's conserved quantities at every gate application, they are unable in principle to approach the known "hydrodynamic" long-time behavior of systems which thermalize. This may be the reason that existing density-matrix methods lose accuracy over time, even though one would expect the accuracy of the matrix-product representation to improve as the state approaches equilibrium.

In this work we propose a truncation of density matrices that ameliorates the positivity problem of Frobenius truncation and exactly preserves the expectation values of all operators on all regions of up to three sites in diameter. (We choose three sites for notational and practical convenience. In Sec. 2.2.4, we discuss how to generalize to regions of arbitrary diameter.)

2.1.2 Intuition: thermalization and computation

Why should one be able to efficiently simulate the dynamics of a local Hamiltonian satisfying the eigenstate thermalization hypothesis? In the long-time limit, the system is well described (as far as local observables are concerned) by a Gibbs state, which has an efficient matrix product density operator (MPDO) representation [85, 154]. Moreover, the system will (one expects) locally thermalize before it reaches global equilibrium, and indeed after some short local thermalization time t_{therm} expectation values of local operators will be well approximated by the expectation values of those operators in a Gibbs state with spatially varying temperature, chemical potential, etc. Heuristically, one can imagine keeping the state exactly out to the local thermalization time and then truncating to an efficient representation. This would require a maximum bond dimension ca. $(d^2)^{vt_{\text{therm}}}$, where v is some entanglement speed and d is the dimension of the onsite Hilbert space. If vt_{therm} is not too large, this approach itself may be workable.

In practice, however, one will wish to efficiently represent the state even at early and intermediate times $t < t_{\text{therm}}$ —and also to avoid dependence on the hard-to-define and likely-unknown constant t_{therm} . Having decided upon an MPDO representation, then, one is faced with the problem of writing a truncation algorithm: an algorithm that will approximate a given MPDO by another, more compact MPDO.

The natural approach, by analogy with matrix product states, is to discard low-weight Schmidt vectors. (This approach turns out to be an imperfect solution, but it is a useful first step.) We call this truncation “Frobenius truncation.” A density operator is a vector in a space with the same tensor-product structure as a state, but a larger onsite dimension. We can therefore cut the chain at bond j into two sections L and R , Schmidt decompose, and truncate it:

$$\rho = \sum_{\alpha=0}^{\chi-1} \hat{x}_{L\alpha} s_{\alpha} \hat{x}_{R\alpha} \mapsto \sum_{\alpha=0}^{\chi'-1} \hat{x}_{L\alpha} s_{\alpha} \hat{x}_{R\alpha}, \quad \chi' < \chi, \quad (2.1)$$

where $\hat{x}_{L\alpha}, \hat{x}_{R\alpha}$ are operators supported on L and R , respectively, and $\text{tr } x_{L\alpha}^{\dagger} x_{L\beta} = \text{tr } x_{R\alpha}^{\dagger} x_{R\beta} = \delta_{\alpha\beta}$. Explicitly, one starts with a matrix-product representation

of the density operator

$$\rho = \sum_{\alpha=0}^{\chi-1} [A_1^{\mu_1} \cdots A_j^{\mu_j}]_{\alpha} s_{\alpha} [B_{j+1}^{\mu_{j+1}} \cdots B_L^{\mu_L}]_{\alpha} \hat{\sigma}_1^{\mu_1} \cdots \hat{\sigma}_L^{\mu_L} \quad (2.2)$$

where $A_l^{\mu_l}, B_l^{\mu_l}$ are $\chi \times \chi$ matrices—with the exception of $A_1^{\mu_1}$ and $B_L^{\mu_L}$, which are $1 \times \chi$ and $\chi \times 1$ respectively. We suppress for compactness a sum on μ (that is, on a basis for the space on onsite operators, here the Pauli matrices σ^{μ} with $\sigma^0 = I$). The truncation (2.1) is then

$$\begin{aligned} \rho &= \sum_{\alpha=0}^{\chi-1} [\cdots A_j^{\mu_j}]_{\alpha} s_{\alpha} [B_{j+1}^{\mu_{j+1}} \cdots]_{\alpha} \hat{\sigma}_1^{\mu_1} \cdots \hat{\sigma}_L^{\mu_L} \\ &\mapsto \sum_{\alpha=0}^{\chi'-1} [\cdots A_j^{\mu_j}]_{\alpha} s_{\alpha} [B_{j+1}^{\mu_{j+1}} \cdots]_{\alpha} \hat{\sigma}_1^{\mu_1} \cdots \hat{\sigma}_L^{\mu_L}, \end{aligned} \quad (2.3)$$

with $\chi' < \chi$. This approximation minimizes the Frobenius (Hilbert-Schmidt) norm distance—but a priori that is not the only norm one could use.

The Frobenius approximation scheme gives poor results for initial states far from equilibrium. One can see why by considering the expectation values of the operator $O_{lt}^y = U(t)\sigma_l^y U(t)^\dagger$ for a system that starts at $t = 0$ in a product of σ^y eigenstates. (Note that O_{lt}^y is a Schrödinger-picture operator parametrized by time t . We work in the Schrödinger picture throughout, except where noted.) At time t , $\langle \psi(t) | O_{lt}^y | \psi(t) \rangle = \pm 1$ —but generically O_{lt}^y will be a large, complicated operator (if we choose t larger than the whole-system entanglement time, as we are free to do, it will have support throughout the system) and essentially unrelated to the local operators we wish to measure. There are 2^L such operators $O_{l_1 l_2 \dots t}^y = U(t)\sigma_{l_1}^y \sigma_{l_2}^y \cdots U(-t)$, all corresponding to long-range operators with expectation value ± 1 . These operators $O_{l_1 l_2 \dots t}^y$ form part of an orthonormal basis for the space of operator space. Errors along the dimensions $O_{l_1 l_2 \dots}^y$ will be penalized by the Frobenius-norm notion of distance with precisely the same severity as errors along more physically relevant dimensions, like σ_l^y . A more reasonable metric for truncation error should be willing to “forget” this information, in favor of more accurately keeping local operators, once they are no longer expected to feedback into the hydrodynamics. (There are more worrying problems still with the naïve Frobenius truncation, which rapidly leads to a dramatic loss of positivity in the supposed density matrix for many far-from-equilibrium initial-conditions,

but these problems can be remedied by considering purifications.)

A BBGKY-like hierarchy for the dynamics of reduced density matrix of a spin chain offers a clue as to how to proceed. In a system governed by a Hamiltonian that is the sum of two-site terms $H = \sum_j h_{j,j+1}$, the dynamics of one-site reduced density matrices (say ρ_j , the reduced density matrix on site j) depends on the two site density matrices:

$$\begin{aligned} \frac{d}{dt}\rho_j &= -i \operatorname{tr}_{\{j' \neq j\}}[H, \rho] \\ &= -i[h_{j,j+1}\rho_{j,j+1}] - i[h_{j-1,j}\rho_{j-1,j}], \end{aligned} \quad (2.4)$$

where we write $\rho_{j,j+1}$ for the two-site reduced density matrix on sites $j, j+1$. Meanwhile the two-site reduced density matrices depend on three-site density matrices, the three-site on four-site, and so on up the sequence. One can imagine truncating this hierarchy at some length l —that is, tracking l -site reduced density matrices for some l , say 2 or 3 or 6, and writing the dynamics for the l -site density matrices in terms of some approximation for the $l+1$ -site density matrices. A natural choice for such an approximation is to replace the $l+1$ -site matrices by their disconnected component, e.g., for $l=1$, take $\rho_{j,j+1} \sim \rho_j \rho_{j+1}$. The problem of the operators $O_{l,t}^y$ then never arises.

The BBGKY-like approach fixes certain problems with Frobenius truncation, but comes with its own set of problems. It is not obvious that the l -site density matrices will even be consistent: it may not be possible to write them as reduced density matrices for a density matrix on the whole system. (Checking this, for a given set of reduced density matrices, is QMA-hard [139].) Moreover, the hard truncation at l sites may not be appropriate to capture the dynamics of the system. Longer-range operators may feed back into the dynamics of few-site density matrices via the hierarchy starting with (2.4)—and conversely, some short-range operators may have a negligible effect on the dynamics of operators of physical interest.

The Frobenius norm attempts to keep all operators with equal weight, while BBGKY keeps connected components only up to a hard cutoff. A natural compromise is to interpolate between the two by weighting the connected components of an operator according to some measure of locality. In the current work, we take the first step in this direction: we approximate the whole-system density matrices in such a way that the dynamics of reduced density matrices on up to three sites matches the BBGKY hierarchy, but

instead of straightforwardly closing the BBGKY hierarchy at that level, we approximate larger connected components using a method similar in spirit to the Frobenius truncation (2.1). Our method zeros out long-range correlations, replacing entanglement between different parts of the system by entanglement with a notional bath.

Although we have used thermalizing Hamiltonians to motivate our method, our method does not in fact assume that the Hamiltonian governing the dynamics is thermalizing. We expect to be able to use it to treat MBL Hamiltonians, with a more stringent accuracy vs. bond dimension tradeoff.

2.2 Method: truncation of MPDOs

Given an MPDO and a particular bond j , we wish to truncate the rank of the MPDO. How can we modify the Frobenius truncation (2.1) in such a way that it does not change the trace of the density matrix, nor the expectation values of local operators? The trick is to start by Schmidt decomposing the whole density matrix, and then cleverly choose basis changes on the spaces of left and right Schmidt vectors that put the data we want to avoid changing in specific, easily-understood locations (see Fig. 2.2). We can then slot our new truncation into a (slight modification of a) standard time-evolution framework, TEBD (see Sec.1.3).

Our truncation algorithm guarantees that the following will not change, up to the precision of the numerical linear algebra involved, in a truncation on bond j (cf Fig. 2.3):

1. the trace of the density matrix, $\text{tr } \rho$;
2. the reduced density matrix $\rho_{1\dots j+1}$ on sites $1, \dots, (j + 1)$; and
3. the reduced density matrix $\rho_{j\dots L}$ on sites j, \dots, L .

Consequently, no truncation will change the expectation of any operator on three contiguous sites, because any such operator is always contained within one of the guaranteed-preserved density operators.

These guarantees do not fully specify our truncation method. To do so, define a matrix of connected correlators across the cut j :

$$\tilde{M}_{\alpha\beta} = \langle \hat{y}_{L\alpha} \hat{y}_{R\beta} \rangle - \langle \hat{y}_{L\alpha} \rangle \langle \hat{y}_{R\beta} \rangle \quad (2.5)$$

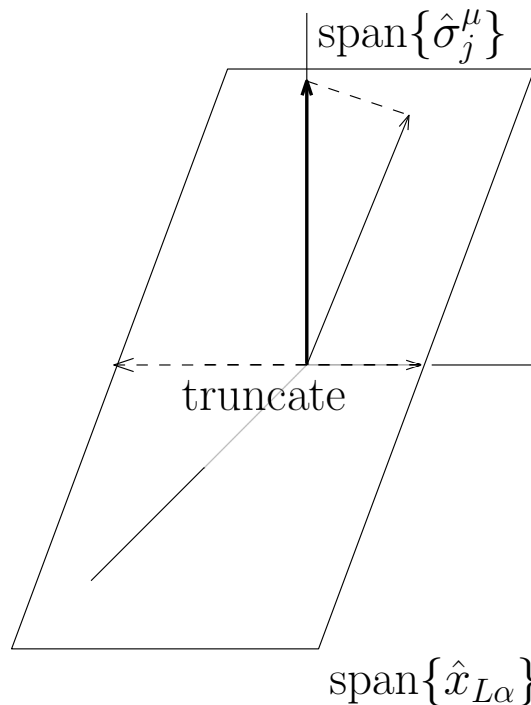


Figure 2.2: DMT, viewed in the space of operators on the left half of the chain. We truncate perpendicular to certain physically-relevant operators (the σ_j^μ).

with $\hat{y}_{L\alpha}$ and $\hat{y}_{R\beta}$ operators supported on sites $1, \dots, j$ and $j+1, \dots, L$ respectively. (The \hat{y} 's spans the set of observables and are defined below.) We wish to replace this matrix \tilde{M} by another \tilde{M}' with lower rank such that $\text{tr}(\tilde{M} - \tilde{M}')^\dagger(\tilde{M} - \tilde{M}')$ is minimized subject to the constraints above. The focus on connected components is itself an important improvement.

2.2.1 Setting, notation, and tools

The concept of the method may be straightforward, but it is obscured by a flurry of notation. We start truncation on bond j with an MPDO of the form


$$\rho = \sum_{\alpha=0}^{\chi-1} \sum_{\{\mu\}} [A_1^{\mu_1} \cdots A_j^{\mu_j}]_\alpha s_\alpha [B_{j+1}^{\mu_{j+1}} \cdots B_L^{\mu_L}]_\alpha \hat{\sigma}_1^{\mu_1} \cdots \hat{\sigma}_L^{\mu_L} \quad (2.6)$$

on an L -site chain. The $A_t^{\mu_t}, B_t^{\mu_t}$ are $\chi \times \chi$ matrices—with the exception of $A_1^{\mu_1}$ and $B_L^{\mu_L}$, which are $1 \times \chi$ and $\chi \times 1$ respectively. (χ , called the bond dimension, will, in fact, vary between bonds and between steps of time-evolution and truncation, but for the moment we suppress this variation.)

In writing our truncation algorithm, we hat our operators ¹. We use Roman

¹We hat all operators except for the density matrix ρ , and identity I .

Truncation at one bond preserves



Truncation at each bond in succession preserves




Figure 2.3: The reduced density matrices that are guaranteed to be preserved under truncation.

letters (frequently j, l) to index sites and bonds; a bond inherits the index of the site to its left. We use Greek letters (frequently α, β, γ —but excepting μ and χ) for the virtual index labeling the Schmidt vectors. The Greek letter $\mu = 0, 1, 2, 3$ is used to label Pauli matrices, i.e., $\hat{\sigma}_j^\mu$ is an operator at site j (with $\hat{\sigma}^0 = I, \hat{\sigma}^1 = \hat{\sigma}^x$, etc.).

Following the standard notation [187], the MPDO is in mixed-canonical form with an orthogonality center at site j —that is, for any $j_1 \leq j$ and $j_2 \geq j$, the operators

$$\begin{aligned} \hat{x}_{L\alpha}[j_1] &= \sum_{\{\mu\}} [A_1^{\mu_1} \cdots A_{j_1}^{\mu_{j_1}}]_\alpha \hat{\sigma}_1^{\mu_1} \cdots \hat{\sigma}_{j_1}^{\mu_{j_1}}, \\ \hat{x}_{R\alpha}[j_2] &= \sum_{\{\mu\}} [B_{j_2+1}^{\mu_{j_2+1}} \cdots B_L^{\mu_L}]_\alpha \hat{\sigma}_{j_2+1}^{\mu_{j_2+1}} \cdots \hat{\sigma}_L^{\mu_L} \end{aligned} \quad (2.7)$$

are orthogonal with respect to the Frobenius inner product

$$\text{tr}[x_{L\alpha}^\dagger x_{L\beta}] = \text{tr}[x_{R\alpha}^\dagger x_{R\beta}] = \delta_{\alpha\beta}.$$

This mixed-canonical form gives the Schmidt decomposition of the density matrix ρ at bond j :

$$\rho = \sum_{\alpha=0}^{\chi-1} \hat{x}_{L\alpha}[j] s_\alpha \hat{x}_{R\alpha}[j]. \quad (2.8)$$

Henceforth, we will implicitly always be working with Schmidt vectors at bond j , and drop the bond label as follows:

$$\begin{aligned}\hat{x}_{L\alpha} &= \hat{x}_{L\alpha}[j], \\ \hat{x}_{R\alpha} &= \hat{x}_{R\alpha}[j].\end{aligned}\tag{2.9}$$

The two vector spaces $\text{span}\{\hat{x}_{L\alpha}\}$ and $\text{span}\{\hat{x}_{R\alpha}\}$ are the setting in which we work.

We frequently abuse notation by replacing s for a diagonal matrix whose entries are s_α . This allows us to shorten Eq. (2.8) into

$$\rho = \hat{x}_L s \hat{x}_R,\tag{2.10}$$

where $\hat{x}_{L,R}$ are row and column vectors of operators, respectively.

With stage set and notation defined, we can walk through the algorithm.

2.2.2 Re-writing the MPDO to expose properties whose preservation we guarantee

We wish to take the MPDO

$$\rho = \sum_{\alpha=0}^{\chi-1} \hat{x}_{L\alpha} s_\alpha \hat{x}_{R\alpha}\tag{2.11}$$

(cut along bond j) and re-write it as

$$\rho = \sum_{\alpha,\beta=0}^{\chi-1} \hat{y}_{L\alpha} M_{\alpha\beta} \hat{y}_{R\beta},\tag{2.12}$$

with the new bases $\{\hat{y}_{L\alpha}\}, \{\hat{y}_{R\alpha}\}$. The bases $\{\hat{y}_{L\alpha}\}, \{\hat{y}_{R\alpha}\}$ and M are chosen such that the properties we wish to avoid changing are characterized by certain easily-identifiable blocks of M :

1. $\text{tr } \rho$ is independent of $M_{\alpha\beta}$ for $\alpha, \beta \neq 0$.
2. The reduced density matrix on sites $1, \dots, (j+1)$, $\rho_{1\dots j+1} = \text{tr}_{\{(j+2)\dots L\}} \rho$, is independent of $M_{\alpha\beta}$ for $\beta \geq 4$.
3. The reduced density matrix on sites j, \dots, L , $\rho_{j\dots L} = \text{tr}_{\{1\dots(j-1)\}} \rho$, is independent of $M_{\alpha\beta}$ for $\alpha \geq 4$.

Once we have made this basis change we will be able to modify $M_{\alpha\beta}$, $\alpha, \beta \geq 4$ with impunity: no such modification will violate our guarantees.

Consider a change of basis

$$\begin{aligned}
\hat{y}_{L\beta} &\equiv (\hat{x}_L Q_L^*)_\beta \\
&\equiv \sum_{\alpha=0}^{\chi-1} \hat{x}_{L\alpha} Q_{L\alpha\beta}^* \\
&= \sum_{\alpha, \{\mu\}} [A_1^{\mu_1} \dots A_j^{\mu_j}]_\alpha Q_{L\alpha\beta}^* \hat{\sigma}_1^{\mu_1} \dots \hat{\sigma}_j^{\mu_j}, \\
\hat{y}_{R\beta} &\equiv (Q_R^\dagger \hat{x}_R)_\beta \\
&\equiv \sum_{\alpha=0}^{\chi-1} \hat{Q}_{R\alpha\beta}^* x_{R\alpha} \\
&= \sum_{\alpha, \{\mu\}} Q_{R\alpha\beta}^* [B_{j+1}^{\mu_{j+1}} \dots B_L^{\mu_L}]_\alpha \hat{\sigma}_{j+1}^{\mu_{j+1}} \dots \hat{\sigma}_L^{\mu_L}
\end{aligned} \tag{2.13}$$

with $Q_{L,R}$ unitary $\chi \times \chi$ matrices. Now write

$$\rho = \hat{x}_L s \hat{x}_R = [\hat{x}_L Q_L^*] [Q_L^T s Q_R] [Q_R^\dagger \hat{x}_R] = \hat{y}_L M \hat{y}_R; \tag{2.14}$$

we can see that M is related to s via

$$M = Q_L^T s Q_R. \tag{2.15}$$

The requisite basis transformations $Q_{L,R}$ are given by QR decompositions

$$\begin{aligned}
Q_{L\alpha\beta} R_{L\beta}^\mu &= \text{tr}[\hat{x}_{L\alpha} \hat{\sigma}_j^\mu] \\
&\propto [A_1^0 \dots A_{j-1}^0 A_j^\mu]_\alpha \in \mathbf{C}^{\chi \times 4}, \\
Q_{R\alpha\beta} R_{R\beta}^\mu &= \text{tr}[\hat{x}_{R\alpha} \hat{\sigma}_{j+1}^\mu] \\
&\propto [B_{j+1}^\mu \dots B_{L-1}^0 B_L^0]_\alpha \in \mathbf{C}^{\chi \times 4}
\end{aligned} \tag{2.16}$$

(here we use the Einstein summation convention). In this context, the fact that the fact that the $R_{L\beta}^\mu$ is upper triangular is exactly the statement that $R_{L\beta}^\mu = 0$ for $\beta > \mu$. (Similarly for $R_{R\beta}^\mu$.)

To see that this is in fact the basis change we seek, first note the trace relations

$$\begin{aligned}\mathrm{tr}[\hat{\sigma}_j^\mu \hat{y}_{L\beta}] &= \sum_{\alpha} \mathrm{tr}[\hat{\sigma}_j^\mu \hat{x}_{L\alpha}] Q_{L\alpha\beta}^* = R_{L\beta}^\mu, \\ \mathrm{tr}[\hat{\sigma}_{j+1}^\mu \hat{y}_{R\beta}] &= \sum_{\alpha} Q_{R\beta\alpha}^\dagger \mathrm{tr}[\hat{\sigma}_{j+1}^\mu \hat{x}_{L\alpha}] = R_{R\beta}^\mu.\end{aligned}\tag{2.17}$$

The trace of the density matrix is

$$\mathrm{tr} \rho = \sum_{\alpha,\beta=0}^{\chi-1} (\mathrm{tr} \hat{y}_{L\alpha}) M_{\alpha\beta} (\mathrm{tr} \hat{y}_{R\beta}) = R_{L0}^0 M_{00} R_{R0}^0.\tag{2.18}$$

(Recall that $\hat{\sigma}_j^0 = \hat{\sigma}_{j+1}^0 = I$ and that the R 's are upper triangular.) This shows that $\mathrm{tr} \rho$ is independent most of the elements of M , as desired. Similarly, the density matrices on sites $1, \dots, j+1$ and j, \dots, L are

$$\begin{aligned}\rho_{1\dots j+1} &= \sum_{\alpha,\beta=0}^{\chi-1} \hat{y}_{L\alpha} M_{\alpha\beta} \sum_{\mu=0}^3 \frac{\mathrm{tr}[\hat{y}_{R\beta} \hat{\sigma}_{j+1}^\mu] \hat{\sigma}_{j+1}^\mu}{2} \\ &= \frac{1}{2} \sum_{\mu=0}^3 \hat{\sigma}_{j+1}^\mu \sum_{\alpha,\beta=0}^{\chi-1} \hat{y}_{L\alpha} M_{\alpha\beta} R_{R\beta}^\mu,\end{aligned}\tag{2.19a}$$

and

$$\begin{aligned}\rho_{j\dots L} &= \sum_{\alpha,\beta=0}^{\chi-1} \sum_{\mu=0}^3 \frac{\hat{\sigma}_j^\mu \mathrm{tr}[\hat{\sigma}_j^\mu \hat{y}_{L\alpha}]}{2} M_{\alpha\beta} \hat{y}_{R\beta} \\ &= \frac{1}{2} \sum_{\mu=0}^3 \hat{\sigma}_j^\mu \sum_{\alpha,\beta=0}^{\chi-1} R_{L\alpha}^\mu M_{\alpha\beta} \hat{y}_{R\beta}.\end{aligned}\tag{2.19b}$$

That is, they depend only on the left four columns and top four rows of M , again as desired.

2.2.3 Modifying the MPDO

Working in the $\{y\}$ bases, we can modify $M_{\alpha\beta}$ for $\alpha, \beta \geq 4$ at will without violating our guarantees. We wish to do so in a way that reduces the rank of M while doing the least violence, in some sense, to the connected components of correlations across the cut. Explicitly, we wish to change the quantities $C(A_L, B_R) = \langle \hat{A}_L \hat{B}_R \rangle - \langle \hat{A}_L \rangle \langle \hat{B}_R \rangle$, where \hat{A}_L, \hat{B}_R have support on the left, right portions of the chain respectively, as little as possible.

First, let us see what the correlator involves. Using the definition $\langle \hat{A} \rangle =$

$\frac{1}{Z} \text{tr} \hat{A} \rho$, with $Z \equiv \text{tr} \rho = R_{L0}^0 R_{R0}^0 M_{00}$, it is easy to see that

$$\langle \hat{A}_L \rangle = \frac{1}{Z} \text{tr}[\hat{y}_{L\alpha} \hat{A}_L] M_{\alpha\beta} R_{R\beta}^0 = \frac{1}{Z} a_\alpha M_{\alpha 0} R_{R0}^0 \quad (2.20a)$$

and

$$\langle \hat{B}_R \rangle = \frac{1}{Z} R_{R\alpha}^0 M_{\alpha\beta} \text{tr}[\hat{y}_{R\beta} \hat{B}_R] = \frac{1}{Z} R_{L0}^0 M_{0\beta} b_\beta. \quad (2.20b)$$

We define $a_\alpha = \text{tr}[\hat{y}_{L\alpha} \hat{A}_L]$, $b_\alpha = \text{tr}[\hat{y}_{R\alpha} \hat{B}_R]$ for convenience. (Throughout this subsection, we employ Einstein summation notation over the Greek indices.)

The expectation value of the product is then

$$\begin{aligned} \langle \hat{A}_L \hat{B}_R \rangle &= \frac{1}{Z} \text{tr}[\hat{y}_{L\alpha} \hat{A}_L] M_{\alpha\beta} \text{tr}[\hat{y}_{R\beta} \hat{B}_R] \\ &= \frac{1}{Z} a_\alpha M_{\alpha\beta} b_\beta. \end{aligned} \quad (2.21)$$

Putting all these together we find

$$C(\hat{A}_L, \hat{B}_R) = \frac{1}{Z} a_\alpha \left(M_{\alpha\beta} - \frac{M_{\alpha 0} M_{0\beta}}{M_{00}} \right) b_\beta. \quad (2.22)$$

Since we wish to only alter pieces of the density matrix which affect correlations across the cut at bond j , we will modify the matrix in the parenthesis, and denote it by

$$\tilde{M}_{\alpha\beta} = M_{\alpha\beta} - \frac{M_{\alpha 0} M_{0\beta}}{M_{00}}. \quad (2.23)$$

At this point we have pushed one step further the process of re-writing the density matrix so that its structure explicitly reflects the distinction between information we are willing to change and information we are not willing to change, but we still have not truncated it. To carry out the truncation, perform an SVD on the lower right block of \tilde{M} , writing $\tilde{M}_{\alpha\beta} = \sum_\gamma X_{\alpha\gamma} r_\gamma Y_{\gamma\beta}$ for $\alpha, \beta \geq 4$. Choose an integer χ' (we will see shortly how it relates to the bond dimension of the final truncated MPDO) and insert a projection $P^{(\chi')}$ onto

the largest χ' elements of r to form a new matrix

$$\tilde{M}' = \left[\begin{array}{c|c} & \\ \hline & XrP^{(\chi')}Y \\ \hline & \end{array} \right]. \quad (2.24)$$

\tilde{M}' differs from \tilde{M} only in those elements $\tilde{M}'_{\alpha\beta}$ with $\alpha, \beta \geq 4$ —that is, in those elements that encapsulate correlations with range ≥ 2 . Moreover, only small elements of M are changed, since we take small (connected) correlations and set them identically to zero.

This truncation results in a new matrix $M'_{\alpha\beta}$ to replace $M_{\alpha\beta}$:

$$M'_{\alpha\beta} = \tilde{M}'_{\alpha\beta} + \frac{M_{\alpha 0}M_{0\beta}}{M_{00}} \quad (2.25)$$

and then

$$\rho' = \hat{y}_{L\alpha} M'_{\alpha\beta} \hat{y}_{R\beta}, \quad (2.26)$$

where the matrix M' has rank at most $8 + \chi'$ (vide infra).

Since we know M' and (matrix-product representations of) the $x_{L\alpha}$ and $x_{R\beta}$, putting this into MPDO form like (2.6) with bond dimension $8 + \chi'$ is a matter of rearrangement. To rearrange into MPDO form perform a second singular value decomposition, this time on M' , for

$$M' = Us'V. \quad (2.27)$$

Since M' has rank at most $\chi' + 8$, there will be at most $\chi' + 8$ nonzero singular values s . The density matrix after truncation is then

$$\rho \mapsto \sum_{\{\mu\}} [A_1^{\mu_1} \cdots A_j^{\mu_j}] s' [B_{j+1}^{\mu_{j+1}} \cdots B_L^{\mu_L}] \hat{\sigma}_1^{\mu_1} \cdots \hat{\sigma}_L^{\mu_L}, \quad (2.28)$$

with

$$\begin{aligned} A_j^{\mu_j} &= A_j Q_L^* U, \\ B_{j+1}^{\mu_{j+1}} &= V Q_L^* B_{j+1}; \end{aligned} \quad (2.29)$$

and the rest of matrices $A_1, \dots, A_{j-1}, B_{j+2}, \dots, B_L$ are untouched. Note that, regardless of our choice of χ' , the reduced density matrices on sites $1, \dots, j+1$,

j, \dots, L are exactly as they were before the truncation.

One loose end needs to be tied up. We claimed that the rank of the matrix M' (hence the bond dimension of the MPDO after truncation) is $\chi' + 8$. To see this, first decompose \tilde{M}' as

$$\tilde{M}' = \tilde{M}'^A + \tilde{M}'^B + \tilde{M}'^C \quad (2.30)$$

with $\tilde{M}'^{A,B,C}$ the left, upper, and lower right blocks of M' respectively:

$$\begin{aligned} \tilde{M}'_{\alpha\beta}{}^A &= \tilde{M}'_{\alpha\beta}, 0 \leq \alpha \leq 3 \\ \tilde{M}'_{\alpha\beta}{}^B &= \tilde{M}'_{\alpha\beta}, 3 < \alpha, 0 \leq \beta \leq 3 \\ \tilde{M}'_{\alpha\beta}{}^C &= \tilde{M}'_{\alpha\beta}, 3 < \alpha, \beta \end{aligned} \quad (2.31)$$

(other elements zero). These have ranks

$$\begin{aligned} \text{rank } \tilde{M}'^A &\leq 4, \\ \text{rank } \tilde{M}'^B &\leq 4, \\ \text{rank } \tilde{M}'^C &\leq \chi', \end{aligned} \quad (2.32)$$

so

$$\begin{aligned} \text{rank}(\tilde{M}') &\leq \text{rank } \tilde{M}'^A + \text{rank } \tilde{M}'^B + \text{rank } \tilde{M}'^C \\ &\leq 8 + \chi'. \end{aligned} \quad (2.33)$$

Since

$$\text{range } \frac{M_{\alpha 0} M_{0\beta}}{M_{00}} \subseteq \text{range } \tilde{M}'^A \quad (2.34)$$

(the range of an operator is also known as its column space) we have

$$\text{rank } M' = \text{rank } \tilde{M}' \leq \chi' + 8, \quad (2.35)$$

as desired.

2.2.4 Preservation of l -site operators

The² preservation of l -site operators for any l is a generalization of this preservation of 3-site operators. Suppose we truncate at the bond between site i and $i + 1$. Then the reduced density matrix of the whole system can once again be

²This section was work with Bingtian Ye.

written

$$\rho = \sum_{\alpha=0}^{\chi-1} \hat{x}_{L\alpha} s_{\alpha} \hat{x}_{R\alpha} \quad (2.36)$$

with

$$\begin{aligned} \hat{x}_{L\alpha} &= \sum_{\{\mu\}} [A_1^{\mu_1} \dots A_i^{\mu_i}]_{\alpha} \hat{\sigma}_1^{\mu_1} \dots \hat{\sigma}_i^{\mu_i} \\ \hat{x}_{R\alpha} &= \sum_{\{\mu\}} [B_{i+1}^{\mu_{i+1}} \dots B_L^{\mu_L}]_{\alpha} \hat{\sigma}_{i+1}^{\mu_{i+1}} \dots \hat{\sigma}_L^{\mu_L}. \end{aligned} \quad (2.37)$$

As before, we perform a basis transformation before SVD decomposition and truncation:

$$\begin{aligned} \hat{y}_{L\beta} &= \sum_{\alpha=0}^{\chi_1} \hat{x}_{L\alpha} Q_{L\alpha\beta}^* \\ \hat{y}_{R\beta} &= \sum_{\alpha=0}^{\chi_1} Q_{R\alpha\beta}^* \hat{x}_{R\alpha}. \end{aligned} \quad (2.38)$$

But now the transformations $Q_{L,R}$ are

$$\begin{aligned} Q_{L\alpha\beta} R_{L\beta}^{\lambda} &= \text{tr}[\hat{x}_{L\alpha} \hat{O}_{i+1-n,i}^{\lambda}] \in \mathbf{C}^{\chi \times 4^n} \\ Q_{R\alpha\beta} R_{R\beta}^{\lambda} &= \text{tr}[\hat{x}_{R\alpha} \hat{O}_{i+1,i+n}^{\lambda}] \in \mathbf{C}^{\chi \times 4^n}, \end{aligned} \quad (2.39)$$

where n is an integer to control the size of the preserved operators, and the $\hat{O}_{j,k}^{\lambda}$ form a basis for operators on the subsystem $[i+1-n, i]$ indexed by λ . After the transformation, we proceed by analogy with the previous section. This procedure will preserve the reduced density matrices of the subsystem $[1, i+n]$ and the subsystem $[i+1-n, L]$. To guarantee this requires the bond dimension $\chi \geq \chi^{\text{preserve}} = 2 \times 4^n$, where 4^n is the number of all possible $\hat{O}_{i+1-n,i}^{\lambda}$ (or $\hat{O}_{i+1,i+n}^{\lambda}$), i.e. the number of operators living in the subsystem $[1, i+n]$ (or the subsystem $[i+1-n, L]$).

Suppose we keep in this way n sites on either side of the cut. What is the largest-diameter operator that is preserved? For an operator on ℓ consecutive sites, we only need to consider the case when the truncation separates it into two parts. Let ℓ_{left} and ℓ_{right} denote the size of the left and the right parts (relative to the truncation) of this operator respectively. Since $\ell_{\text{left}} + \ell_{\text{right}} = \ell$, we have $\min\{\ell_{\text{left}}, \ell_{\text{right}}\} \leq \lceil \ell/2 \rceil$. If $\ell_{\text{left}} \leq n$, the ℓ -site operator will live on the subsystem $[1, i+n]$, and is thus preserved (the marginal case is the blue frame in Fig. 2.4a). Similarly, if $\ell_{\text{right}} \leq n$, the ℓ -site operator will

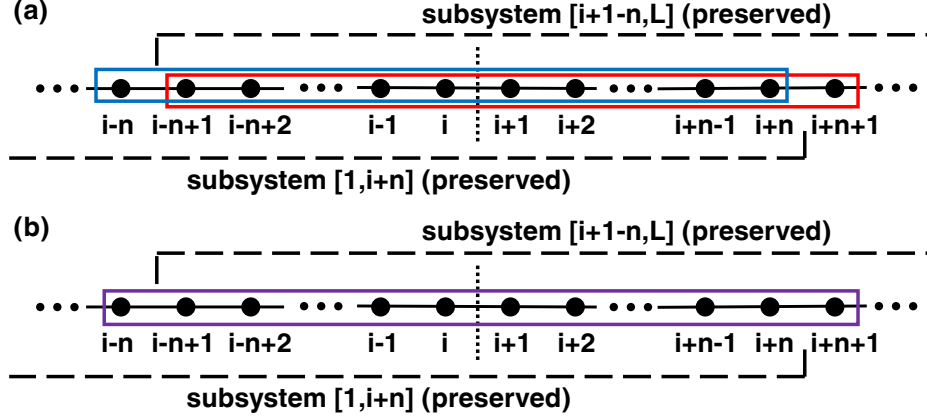


Figure 2.4: The operators preserved during the truncation. (a) $2n + 1$ is the maximum size of the operators that can be preserved. (b) A truncation can change the expectation of a $(2n + 2)$ -site operator.

also be preserved (the marginal case is shown by the red frame in Fig. 2.4a). Therefore, any l -site operator with $\lceil l/2 \rceil \leq n$ is preserved during a truncation on any bond i , which means for a given n , we can preserve all $(2n + 1)$ -site operators. However, a $(2n + 2)$ -site operator can be changed by the truncation at the middle of it (Fig. 2.4b). Combining the previous expression for $\chi^{preserve}$ and $\ell = 2n + 1$, we prove that to preserve all ℓ -site operators requires bond dimension $\chi^{preserve} = 2^\ell$.

2.3 Results

We apply the method to a variety of initial states, pure and mixed. We time evolve by the boustrophedon Trotter decomposition (1.114) of a Hamiltonian known to satisfy the eigenstate thermalization hypothesis with Trotter step $\delta t = 1.0$ (except where specified). We work at a maximum bond dimension cutoff (i.e., at each gate application we truncate to this cutoff using the algorithm described in section 2.2) and measure performance by varying this cutoff.

2.3.1 Hamiltonian

We take as our Hamiltonian the spin-1/2 transverse-field Ising model on an L -site chain with open boundary conditions:

$$H = \sum_{j=1}^{L-1} S_j^z S_{j+1}^z + \frac{1}{2} h^x \sum_{j=1}^L S_j^x + \frac{1}{2} h^z \sum_{j=1}^L S_j^z. \quad (2.40)$$




Figure 2.5: $\langle S^z \rangle$ for the initial state described in (2.42) on a 24-site chain.

At $h^z = 0$ the model is integrable (by Jordan-Wigner and Bogoliubov transformations); the longitudinal field $h^z \sum S^z$ breaks integrability, and at $h^z = 0.8090, h^x = 0.9045$ the model is known to satisfy the eigenstate thermalization hypothesis in a strong sense [115]. We work at onsite fields

$$h^z = 0.8090, \quad h^x = 0.9045 \quad (2.41)$$

(except where otherwise specified). Despite their ETH nature, TFIM Hamiltonians like this can display ill-understood pre-thermalization behavior, thought to be related to long-range emergent conserved or nearly-conserved quantities [13, 113, 135]. We do not expect DMT to be able to capture this emergent-integrable behavior (see Section 2.1), so we choose our initial conditions to avoid it.

2.3.2 Application: pure-state evolution

We engineer an initial state with a long-wavelength energy variation by taking a product of σ^y eigenstates and rotating blocks of four spins alternately towards $+z$ and $-z$ (cf Fig. 2.5). The initial state is

$$|\text{near}Y\rangle = \prod_{j=1}^L [1 + i(1 + g_j)\sigma_j^+] |\downarrow\downarrow \dots \downarrow\rangle \quad (2.42)$$

(suitably normalized), where

$$g_j = 0.1 \times \begin{cases} -1 & j \bmod 8 = 1, 2, 7, \text{ or } 0, \\ +1 & j \bmod 8 = 3, 4, 5, \text{ or } 6. \end{cases} \quad (2.43)$$




Figure 2.6: Normalization $Z = [\text{tr} \rho] / \sqrt{\text{tr}[\rho^2]}$ as a function of time, comparing DMT (solid) and Frobenius (dashed), for the initial pure state (2.42) evolving under the Hamiltonian (2.40). Note that the second Rényi entropy of the whole chain (64-sites long) is $S_2 = 2 \ln Z$. As $Z = 1$ for a pure state, any deviations from $Z = 1$ result from the truncation.

(We choose the state to be near the σ^y product state in order that we may avoid the pre-thermalization behavior found in [13, 113, 135].)

Since the initial state is a product state, it may be represented exactly as an MPO with bond dimension $\chi = 1$. Trotter time evolution increases the bond dimension with each time step δt , but truncation (whatever the algorithm) kicks in only at a time $t_{\text{trunc}}(\chi_{\text{max}})$ when $\chi(t)$ reaches χ_{max} . Thus for each χ_{max} the time evolution is semiexact (that is, exact up to error resulting from the Trotter decomposition of the Hamiltonian) for $t < t_{\text{trunc}}(\chi_{\text{max}})$, at which time it begins to deviate from the semiexact value. This effect appears in all of our results; we also use it to benchmark our method (vide infra).

Figure 2.6 shows the normalization $Z = \frac{\text{tr} \rho}{\sqrt{\text{tr} \rho^2}}$ as a function of time. The normalization is related to the second Rényi entropy of the entire chain $S_2 \equiv -\ln \frac{\text{tr}[\rho^2]}{[\text{tr} \rho]^2}$ via $S_2 = 2 \ln Z$. Time evolution by DMT produces bath entropy for the system, and this is reflected in the increase of Z as a function of time. In contrast, we find that the Frobenius method produces non-physical states with $Z < 1$ over the course of time-evolution, which results from negative eigenvalues of the density matrix generated in the truncation. The observation that $Z \geq 1$ for DMT does not imply positive semi-definiteness, but suggests that any error arising from the negative eigenvalues is small and well-controlled.

Figure 2.7 shows the second Rényi entropy of the left half of the chain—that is, the subsystem consisting of sites 1 to $L/2 = 32$. (The von Neumann entropy is difficult to calculate for MPDOs, while second Rényi entropies are nearly

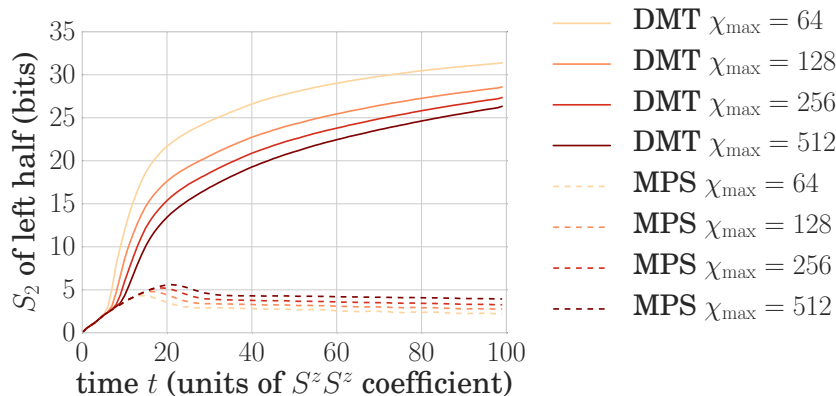


Figure 2.7: Second Rényi entropy (in units of bits) of left half of a 64-site chain for an initial pure state (2.42) evolving under the Hamiltonian (2.40), in matrix product state (MPS) and DMT simulations. The largest entropy we see ($\chi = 64$ at $t = 100$) is $S_2 \approx 31.4$ bits, very close to the theoretical maximum of 32 bits.

trivial.) In contrast to matrix product states, MPDOs can represent states with arbitrarily large entropy by replacing system entanglement entropy with bath entropy.

Note that once truncation starts, the entropy in the DMT simulation increases above that in MPS. This is not unexpected: in ordinary matrix product state TEBD the entanglement entropy of the left half is exactly its entropy with the right half and is a property of the matrix product state at bond $\lfloor L/2 \rfloor$, so it can only increase when we apply a gate at bond $\lfloor L/2 \rfloor$. In the DMT algorithm, on the other hand, the entanglement entropy of the left half of the chain is entanglement entropy not only with the right half but also with a notional bath, and it increases with every truncation on bonds within the left half.

Figure 2.8 shows the system’s total energy over time as simulated by ordinary matrix product state TEBD and our density-matrix TEBD. In the DMT simulation, the energy is constant. Matrix product state time evolution, however, imposes an additional ‘heating’ whereas DMT is designed to conserve total energy. Because the matrix product state representation is biased towards low-entanglement states, the system drifts towards the extrema of the energy spectrum over time. (The simulation begins with positive energy, and hence drifts towards the negative temperature state $T \rightarrow 0^-$.)

Onsite spins are easy to measure but hard to analyze: their expectation values

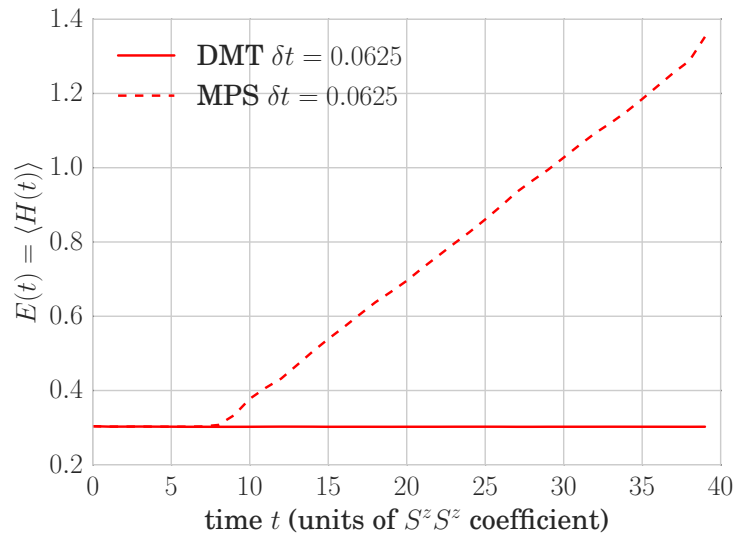


Figure 2.8: Energy over time at fixed $\chi = 16$ for the initial pure state (2.42) evolving under the Hamiltonian (2.40) at system size $L = 64$. By design, the total energy remain constant under DMT.

are noisy and they decay quickly with time. Instead, we measure a Fourier component of the energy density. The energy density ϵ_j , defined over a pair of sites, is

$$\begin{aligned}
 \epsilon_1 &= \frac{h^z}{2} (S_1^z + \frac{1}{2} S_2^z) + \frac{h^x}{2} (S_1^x + \frac{1}{2} S_2^x) + S_1^z S_2^z, \\
 \epsilon_{1 < j < L-1} &= \frac{h^z}{2} (\frac{1}{2} S_j^z + \frac{1}{2} S_{j+1}^z) + (x \leftrightarrow z) + S_j^z S_{j+1}^z, \\
 \epsilon_{L-1} &= \frac{h^z}{2} (\frac{1}{2} S_{L-1}^z + S_L^z) + (x \leftrightarrow z) + S_{L-1}^z S_L^z.
 \end{aligned} \tag{2.44}$$

We measure a Fourier component of the energy density

$$\epsilon_{k=\pi/4} = -\frac{1}{L} \sum_{j=1}^{L-1} e^{ikj} \epsilon_j, \quad k = \pi/4 \tag{2.45}$$

with a wavelength of 8 sites. Fourier components are eigenmodes of the diffusion equation which should govern the system's long-time non-equilibrium behavior. We choose this particular component (and choose the initial state accordingly) because its wavelength is long enough that it should not immediately decay, but not so long as to be longer than accessible system sizes.

Figure 2.9 shows the the Fourier component of energy density (2.45) as simulated by matrix product states (dashed) and DMT (solid) for $L = 16, 20, 24$.




Figure 2.9: (a) Fourier component of energy density for the initial pure state (2.42) evolving under the Hamiltonian (2.40) on chains of length $L = 16, 20, 24$. (b) The ‘error’ of the energy density, measured by comparing each data set with the semiexact result as simulated by matrix product states at $\chi = 2^{L/2}$. Note we do not show the MPS simulations for $\chi_{\max} = 2048, 4096$.

At fixed maximum bond dimension $\chi \ll 2^{\lfloor L/2 \rfloor}$, DMT is more accurate than matrix product state TEBD, which illustrates the power of DMT in both short- and long-time dynamics. Moreover, where matrix product state TEBD error increases with system size, DMT error decreases. This is due to finite-size deviations from thermalizing behavior: oscillations about the local equilibrium values for local operators result from long-range coherences that we do not expect to be able to capture.

Any pure state on a system of length L can be represented exactly by matrix product states with bond dimension $\chi_{\max} = 2^{\lfloor L/2 \rfloor}$. At this bond dimension the evolution by matrix product state TEBD becomes semiexact for all times; no truncation occurs for the simulation. We can therefore simulate pure-state evolution of a matrix product state using exactly the same Hamiltonian (2.40) and boustrophedon Trotter decomposition (1.114), and compare the results to those of DMT, shown in Fig. 2.9. The data is a measure of the error introduced

by the truncation; these are small $\approx 10^{-3}$ for a wide range of bond dimensions in the DMT simulations.

2.3.3 Application: mixed-state evolution (near equilibrium)

To probe the behavior of our algorithm near equilibrium, we take as our initial state a Gibbs state with a spatially-varying temperature

$$\rho_0 \propto \exp \left[\sum_j \beta_j \epsilon_j \right] \quad (2.46)$$

with ϵ_j the energy density of Eq. (2.44), and

$$\beta_j = \beta_0(1 + g'_j) \quad (2.47)$$

where

$$g'_j = 0.1 \times \begin{cases} 0 & j \bmod 8 = 1, 2, 7, \text{ or } 0, \\ 1 & j \bmod 8 = 3, 4, 5, \text{ or } 6. \end{cases} \quad (2.48)$$

This temperature profile is broadly similar to the S^z profile we impose on the pure initial state (see Eq. (2.42) and Fig. 2.5).

See Sec. 1.4.4.1 for details of the construction of the Gibbs state.

In Figs. 2.10 and 2.11 we compare DMT to the purification method of Karasch, Bardarson and Moore [104], which we label “purification”. This method takes advantage of the freedom to apply unitaries to the ancillae by time-evolving the ancillae backwards even as it evolves the physical system forwards. The time-evolution framework is therefore very similar to ours; the chief differences are in the interpretation of the vector space $(\mathbf{C}^{d^2})^L$ in which one works and in the truncation algorithm. The similarity is magnified by our choice to use the boustrophedon Trotter decomposition (1.114) not only for DMT but for purification.

Both DMT and this purification time evolution converge very quickly, as one might expect: the results are essentially identical between the methods and between bond-dimension cutoffs, even to quite small bond dimensions. (Note that we subtract the thermal value in each case.)

2.3.4 Application: mixed-state evolution (far from equilibrium)

One might worry that the two initial states (2.42) from Sec. 2.3.2 and (2.46) from Sec. 2.3.3 are each special cases in their own ways: the first is a pure




Figure 2.10: Fourier component of energy density for (a) purification time evolution and (b) DMT for the near-equilibrium mixed state (2.46) evolving under the Hamiltonian (2.40) on a 128-site chain. The thermal value is $\langle \epsilon_{k=\pi/4} \rangle_{\text{thermal}} = -0.00038$.

state, and the second is very near equilibrium.

In order to probe the performance of DMT for in more generic situations, we quench from a Gibbs state of the TFIM (2.40) with

$$\begin{aligned} h_0^x &= 0.5 \\ h_0^z &= 0.5 \end{aligned} \tag{2.49}$$

and

$$\beta_j = \beta_0(1 + g'_j), \tag{2.50}$$

where

$$g'_j = 0.1 \times \begin{cases} 0 & j \bmod 8 = 1, 2, 7, \text{ or } 0, \\ 1 & j \bmod 8 = 3, 4, 5, \text{ or } 6 \end{cases} \tag{2.51}$$

(as in section 2.3.3), to a TFIM (2.40) with

$$\begin{aligned} h_1^x &= 2.0 \\ h_1^z &= 0.5. \end{aligned} \tag{2.52}$$

We again compare to the purification method [104], and find that our method





Figure 2.11: Expectation value of S^z at the midpoint of the chain for (a) purification time evolution and (b) DMT for the near-equilibrium initial state (2.46) evolving under the Hamiltonian (2.40) on a 128-site chain. Both methods converge very quickly, so they give nearly identical results (cf. Fig. 2.15). This expectation value fails to approach the thermal value due to the large Trotter step we use ($dt = 1.0$). The thermal value is $\langle S^z_{[L/2]} \rangle_{\text{thermal}} = -0.0622$.

and that purification time evolution both converge quickly (see Figs. 2.12a and 2.13). Even very small bond dimensions (e.g. $\chi = 16$) can accurately treat long-time, hydrodynamic behavior; accurately treating short-time behavior requires somewhat higher bond dimension.



(a) Fourier component of energy density for purification time evolution and DMT starting from a far-from-equilibrium initial state on a 128-site chain. The thermal value is $\langle \epsilon_{k=\pi/4} \rangle_{\text{thermal}} = -0.00031$.




Figure 2.13: Expectation value of S^z at the midpoint of the chain for purification time evolution and DMT starting from a far-from-equilibrium initial state on a 128-site chain. The thermal value is $\langle S^z_{[L/2]} \rangle_{\text{thermal}} = -0.021$.

2.3.4.1 Convergence of mixed-state evolution

It is difficult to judge convergence of any of the three algorithms from plots like Figure 2.10 or 2.12a. In Figure 2.14 we take a near-equilibrium initial state and plot the deviation in $\epsilon_{k=\pi/4}$, as measured for a series of bond dimensions χ , from the last (largest) χ in the series. In Figure 2.16 we do the same for a far-from-equilibrium mixed state, and in Figures 2.15, 2.17 for $S^z_{L/2}$ for near-equilibrium and far-from-equilibrium mixed states.

Our method converges with approximately the same bond dimension vs. accuracy tradeoff as purification time evolution for both the near-equilibrium initial state (Fig. 2.14) and the far-from-equilibrium initial state (Fig. 2.16). In both cases, Frobenius time evolution converges more slowly than either method.

2.4 Conclusion

We have presented an algorithm for approximating density operators by low-rank matrix product operators suitable for simulating long-time quantum dynamics. The method exactly preserves expectation values of operators on up to three contiguous sites, and it slots neatly into a standard Trotter-decomposition framework for time evolution of matrix product structures (TEBD), allowing time evolution by an ETH Hamiltonian of a variety of initial states.




Figure 2.14: Convergence of $\epsilon_{k=\pi/4}$ for three algorithms. Initial state is a near-equilibrium mixed state (cf Section 2.3.3 and Figure 2.10) on a 128-site chain. For each algorithm, we plot $\epsilon_{k=\pi/4}[\chi] - \epsilon_{k=\pi/4}[\chi = \chi_{\max}]$ —that is, how far the measurement during a run with a certain bond dimension χ deviates from measurement during a run with some high bond dimension.




Figure 2.15: Convergence of S^z at site $L/2$ for three algorithms. Initial state is a near-equilibrium mixed state (cf Section 2.3.3 and Figure 2.11) on a 128-site chain. For each algorithm, we plot $S^z_{L/2}[\chi] - S^z_{L/2}[\chi = \chi_{\max}]$ —that is, how far the measurement during a run with a certain bond dimension χ deviates from measurement during a run with some high bond dimension.

Our algorithm, DMT, qualitatively outperforms its nearest competitor (ordinary matrix product state TEBD) for pure initial states. We use the fact that matrix product density operators with small bond dimension can represent states with high entropy to circumvent the area-law entanglement bound on matrix product states. Thus far the work is unoriginal: Zwolak and Vidal realized this was possible more than a decade ago. Our key insight is that we can preserve the trace of the density matrix and the expectation values of conserved quantities by appropriately rotating the Schmidt spaces at the bond at which we truncate. Consequently, DMT can simulate time evolution by ETH Hamiltonians to arbitrary times using memory and computation time polynomial in system size.




Figure 2.16: Convergence of $\epsilon_{k=\pi/4}$ for three algorithms. Initial state is a far-from-equilibrium mixed state (cf Section 2.3.4 and Figure 2.12a) on a 128-site chain. For each algorithm, we plot $\epsilon_{k=\pi/4}[\chi] - \epsilon_{k=\pi/4}[\chi = \chi_{\max}]$ —that is, how far the measurement during a run with a certain bond dimension χ deviates from measurement during a run with some high bond dimension.




Figure 2.17: Convergence of S^z at site $L/2$ for three algorithms. Initial state is a far-from-equilibrium mixed state (cf Section 2.3.4 and Figure 2.13) on a 128-site chain. For each algorithm, we plot $S_{L/2}^z[\chi] - S_{L/2}^z[\chi = \chi_{\max}]$ —that is, how far the measurement during a run with a certain bond dimension χ deviates from measurement during a run with some high bond dimension.

In addition, DMT matches the current state of the art (purification time evolution) for near-equilibrium mixed initial states and outperforms it for far-from-equilibrium initial states.

The reader would be right to worry that our method does not converge: as we increase the bond dimension above a certain value (perhaps 2^5 – 2^6), the accuracy of our method does not improve. We suspect that—once again—this is a result of the operators $O_{i,t}^y$, whose large expectation values result from the fact that we start near an S_i^y eigenstate. When we reduce the rank of the matrix \tilde{M} in (2.24), we still do so in a way that minimizes error with respect to the Frobenius norm (even though we have arranged to exactly preserve very-short-range operators). This means that the operators $O_{i,t}^y$ again dominate

the error, and the matrix resulting from the truncation is pulled toward those operators. The obvious next step is to reduce rank in such a way that we minimize error with respect to a different norm, one that takes into account the spatial structure of the operator space: if we truncate at bond j , we should weight errors along $\sigma_{j-1}^z \sigma_{j+2}^z$ more heavily than errors along $\sigma_{j-7}^z \sigma_{j+6}^z$. Such controlled-metric truncation is a natural extension of this work.

One natural question to attack using our algorithm is the characterization of the ergodic side of the MBL transition. The random field Heisenberg model with small disorder appears to satisfy the ETH [160], but the nature of its dynamics is unclear (see the review of Luitz and Bar Lev [143]). Quantities like the spin-spin correlation $\langle S_{i+r}^z(t) S_i^z(t) \rangle$, from which one can compute a number of diagnostics for subdiffusion, should be straightforward to calculate using our method.

More interesting still are questions about interfaces between ETH and MBL systems. Besides being of inherent interest (how large must a bath be to thermalize an MBL system of a given size? how quickly does it thermalize?), answers to these questions will shed light on the phenomenological RG schemes of Potter, Vasseur, and Parameswaran [172] and Vosk, Huse, and Altman [220] for which ETH-MBL interfaces are fundamental building blocks. Because MBL systems display low entanglement in a wide variety of situations, we expect our algorithm to be able to simulate both bath and system out to large system sizes.

HEATING, PRETHERMALIZATION, AND HYDRODYNAMICS IN A HIGH-FREQUENCY FLOQUET SYSTEM

In the previous chapter we considered out-of equilibrium dynamics resulting from a quench. But, a many-body system can also be taken out of equilibrium via periodic (Floquet) driving—a strategy which has received recent attention in the context of novel Floquet phases of matter [38, 56, 110, 116, 138, 216–218, 235, 243]. In this case, the non-equilibrium system is generically expected to absorb energy from the driving field (so-called Floquet heating) until it approaches a featureless infinite temperature state [29, 45, 128, 150, 173, 174].

Floquet heating, emergent hydrodynamics and microscopic thermalization are all fall naturally under the umbrella of non-equilibrium dynamics, but understanding the interplay between them represents an important step toward the characterization and control of non-equilibrium many-body systems [4, 6, 7, 30, 57, 126, 155, 224]. One can already see such connections; for example, in the limit of a high-frequency Floquet drive, energy absorption is set by an extremely slow heating rate. So one anticipates a relatively long timescale where the system’s stroboscopic dynamics can be captured by an effective static *prethermal Hamiltonian*. These expectations prompt one to ask: How do the late-time dynamics of driven systems account for both the prethermal Hamiltonian’s hydrodynamics and the energy absorption associated with Floquet heating?

In this chapter, we investigate the dynamics of a non-integrable Floquet spin chain using DMT. Our main results are three fold. First, by benchmarking DMT against conventional Krylov subspace methods (for small systems, $L = 20$), we find that DMT efficiently simulates the system’s late-time physics, including both prethermalization and late-time Floquet heating to infinite temperature (Fig. 3.2). Second, by extending to larger systems (up to $L = 100$), we essentially eliminate the finite-size effect and demonstrate the exponentially slow Floquet heating (Fig. 3.1); such a calculation is fundamentally impossible for either exact diagonalization based methods (owing to the size of the Hilbert space) or conventional density matrix renormalization group (DMRG)

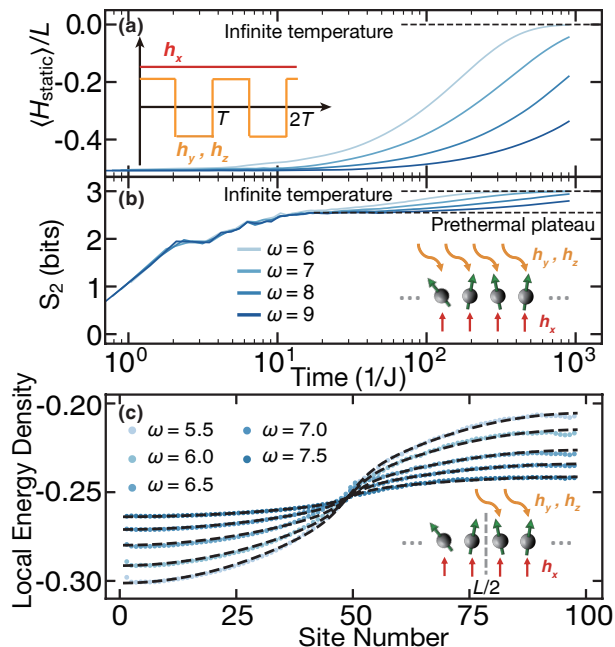


Figure 3.1: Floquet thermalization of an $L = 100$ spin chain ($\chi = 128$). (a) Average energy density measured with respect to $D_{\text{eff}}^{(0)} = H_{\text{static}}$ under a global drive. (b) The second Rényi entropy of the leftmost three sites. (c) Spatial profiles of energy density under a half-system drive with $\langle H_{\text{static}} \rangle = -0.25$. Insets: the drive's time dependence (a) and schematics of the global drive (b) and the half-system drive (c).

methods (owing to the large amount of entanglement at late times). Finally, we illustrate the emergence of a classical hydrodynamical description from the underlying quantum model. Having studied the late-time heating, we begin by considering transport in the associated static model and use DMT to directly measure the diffusion coefficient of the energy density. We then consider an inhomogeneously driven spin chain, where both heating and transport are crucial for understanding the late-time equilibration of the system (Fig. 3.8). Intriguingly, the dynamics of the local energy density is well described by a simple hydrodynamical equation.

3.1 Model and Floquet phenomenology

We study the dynamics of a one-dimensional spin-1/2 chain whose evolution is governed by a time periodic Hamiltonian $H(t) = H_{\text{static}} + H_{\text{drive}}(t)$, where

$$H_{\text{static}} = \sum_{i=1}^{L-1} [J\sigma_i^z \sigma_{i+1}^z + J_x \sigma_i^x \sigma_{i+1}^x] + h_x \sum_{i=1}^L \sigma_i^x, \quad (3.1)$$

with σ_i^α being the $\alpha \in \{x, y, z\}$ Pauli operator acting on site i .¹ We note that, while the bond terms can be mapped to a free-fermion integrable model, the additional field term breaks this integrability. The driven part, $H_{\text{drive}}(t) = H_{\text{drive}}(t + T)$, has a period $T = 2\pi/\omega$ and corresponds to an oscillating field in the \hat{y} and \hat{z} directions (Fig. 3.1a inset):

$$H_{\text{drive}}(t) = \sum_{i=1}^L v_i(t) (h_y \sigma_i^y + h_z \sigma_i^z). \quad (3.2)$$

In this work we consider two driving protocols: a **global drive**, where all the spins are driven [$v_i(t) = \text{sgn} \cos(\omega t)$ for all spins], and a **half-system drive**, where only the right half of the system is driven [$v_{i \leq L/2}(t) = 0$ and $v_{i > L/2}(t) = \text{sgn} \cos(\omega t)$]. In our calculations we take the initial state to be the Néel state with a domain wall every four spins while using parameters $\{J, J_x, h_x, h_y, h_z\} = \{1, 0.75, 0.21, 0.17, 0.13\}$ within the high-frequency regime $\omega \geq 5J$. Based on previous studies we believe this choice of parameters and initial state to be generic and to capture the main features of Floquet heating [150].

The quenched dynamics of a high-frequency driven system is characterized by two timescales. The heating timescale, τ^* , determines the rate of energy absorption from the drive (thus the approach to infinite temperature) and has been proven to be at least exponential in the frequency of the drive, $\tau^* \geq \mathcal{O}(e^{\omega/J_{\text{local}}})$, where J_{local} is a local energy scale [4, 6, 7, 57, 126, 155]. Until τ^* the dynamics of the system is, stroboscopically, well described by the *static* prethermal Hamiltonian $D_{\text{eff}} = H_{\text{static}} + \mathcal{O}(\omega^{-1})$. The prethermalization timescale, $\tau_{D_{\text{eff}}}$, determines the time at which the system approaches the equilibrium state with respect to D_{eff} . When $\tau_{D_{\text{eff}}} \ll \tau^*$, the system exhibits a well defined, long-lived prethermal regime with respect to the prethermal Hamiltonian.

There is some ambiguity in the definition of the heating timescale τ^* . It is the characteristic timescale on which the system's state decays towards the infinite-temperature state. But how does one measure that? Two options present themselves: the energy and the entropy. The energy is perhaps the most physically natural—in an ETH system, one expects that any state with energy $E = 0$ will be locally like the infinite temperature state, or at least

¹A Jordan-Wigner transform maps H_{static} to a $p + ip$ superconductor with hopping $t = J + J_x$, pairing $\Delta = J - J_x$, and an integrability breaking term of order h_x .




Figure 3.2: Comparison of the evolution of a $L = 20$ spin chain under global drive between DMT and Krylov subspace methods. (a) Average energy density $\langle H_{\text{static}} \rangle / L$. (b) A typical local observable $\sigma_9^z \sigma_{10}^z$. (c) The second Rényi entropy S_2 of the leftmost 3 sites. Insets (early-time behaviors): (b) errors in DMT $\delta_{\langle \sigma_9^z \sigma_{10}^z \rangle} = \langle \sigma_9^z \sigma_{10}^z \rangle_{\text{DMT}} - \langle \sigma_9^z \sigma_{10}^z \rangle_{\text{Krylov}}$, (c) early-time evolution of S_2 . The blue lines are $\chi = 64$; the red lines in the insets are $\omega = 10$.

evolve into a state locally like the infinite temperature state on a timescale like $\tau_{D_{\text{eff}}} \ll \tau^*$, so one can measure the τ_E^* such that

$$\langle E(t) \rangle \sim -e^{-t/\tau_E^*}. \quad (3.3)$$

But the entropy, in particular the second Rényi entropy, is more natural from a formal point of view. Take a subsystem of length l and consider the Frobenius distance between the reduced density matrix ρ_l on the subsystem and the infinite-temperature density matrix $2^{-l}I$. This characterizes how far *all* operators are from their infinite temperature values. If we decompose the reduced density matrix into Pauli-string components

$$\rho_l = 2^{-l} \left[I + \sum_{\vec{\nu}} R_{\vec{\nu}} \sigma^{\vec{\nu}} \right], \quad (3.4)$$

where $\vec{\nu} = \nu_1\nu_2\nu_3\dots\nu_l, \nu_j \in 0\dots 3$] label Pauli strings $\sigma^{\vec{\nu}}$ and we take each $\vec{\nu}$ to have at least one $\nu_j \neq 0$, then the $R_{\vec{\nu}}$ are the expectation values of those Pauli strings

$$\text{tr}[\sigma^{\vec{\nu}}\rho_l] = R_{\vec{\nu}} \quad (3.5)$$

and

$$\|\rho_l - I\|^2 = 2^{-2l} \text{tr}(\rho_l - I)^2 = 2^{-l} \sum_{\vec{\nu}} R_{\vec{\nu}}^2 = 2^{-l} \sum_{\vec{\nu}} (\text{tr}[\sigma^{\vec{\nu}}\rho_l])^2. \quad (3.6)$$

But this is intimately related to the second Rényi entropy:

$$\begin{aligned} \|\rho_l - 2^{-l}I\|^2 &= \text{tr}\rho_l^2 + 2^{-2l} \text{tr} I - 2^{-l} \text{tr} \rho_l \\ &= 2^{-S_2} - 2^{-l} \\ &= 2^{-l} (e^{\ln 2(l-S_2)} - 1) \\ &\simeq 2^{-l} \ln 2(l - S_2) \end{aligned} \quad (3.7)$$

when S_2 is close to maximum the entropy l . (We measure the entropy S_2 in bits, not nits). Then the time τ_S^* such that for the subsystem

$$l \ln 2 - S_2(t) \sim e^{-t/\tau_S^*} \quad (3.8)$$

characterizes the decay of *all* observables

But there turns out to be a factor of two between τ_E^* and τ_S^* . This is already apparent from Eq. (3.6): τ_S^* actually characterizes the decay time of the *square* of the local expectation values, including energy. But we can see the factor more explicitly when ρ_l is a high-temperature Gibbs state.

Define $\tau_S^* = \frac{1}{2}\tau_S'^*$, and consider a Gibbs ensemble at temperature T . The probability p_i assigned to the i 'th eigenstate (with ϵ_i being its eigenenergy) of H_{static} can be approximated to the first order as

$$p_i = \frac{e^{-\beta\epsilon_i}}{\text{tr}[e^{-\beta H}]} \approx \frac{1 - \beta\epsilon_i}{\text{tr}[1 - \beta H]} = 2^{-L}(1 - \beta\epsilon_i), \quad (3.9)$$

where $\beta = 1/T$, 2^{-L} is the dimension of the Hilbert space, and we use the fact that $\text{tr} H = 0$. The energy is then

$$E = \sum_i p_i \epsilon_i \approx 2^{-L} \sum_i (1 - \beta\epsilon_i) \epsilon_i = -2^{-L} \sum_i \beta \epsilon_i^2 \propto \beta. \quad (3.10)$$

The second Rényi entropy of the entire system is similar:

$$\begin{aligned}
S_{\text{entire}} &= -\log_2 \sum p_i^2 \\
&\approx -\log_2 \left[2^{-2L} \sum_i (1 - \beta \epsilon_i)^2 \right] \\
&= L - \beta^2 2^{-L} \sum_i \epsilon_i^2
\end{aligned} \tag{3.11}$$

for

$$\Delta S_{\text{entire}} \equiv L - S_{\text{entire}} \simeq \beta^2 2^{-L} \sum_i \epsilon_i^2 \propto \beta^2. \tag{3.12}$$

Since the entropy is an extensive quantity for a Gibbs state, one expects this behavior to hold for any subsystem; thus $\Delta S \propto \beta^2$. Therefore

$$\begin{aligned}
\frac{1/\tau_E^*}{1/\tau_S^*} &= \left(\frac{1}{E} \frac{dE}{dt} \right) / \left(\frac{1}{\Delta S} \frac{d\Delta S}{2dt} \right) \\
&= \left(\frac{d \log |E|}{dt} \right) / \left(\frac{d \log \Delta S}{2dt} \right) \\
&= \left(\frac{d \log |E|}{d \log \beta} \right) / \left(\frac{1}{2} \frac{d \log \Delta S}{d \log \beta} \right) \\
&= 1.
\end{aligned} \tag{3.13}$$

In Figs. 3.1a,b, we illustrate these two timescales τ^* and $\tau_{D_{\text{eff}}}$ for an $L = 100$ initial state evolved with DMT (we defer the analysis of the validity of DMT until next section). In Fig. 3.1a, we highlight τ^* using the average energy density $\langle H_{\text{static}}(t) \rangle / L$. Until τ^* , the evolution is well described by D_{eff} , and thus $\langle H_{\text{static}} \rangle$ remains constant (up to ω^{-1} corrections). After τ^* , $\langle H_{\text{static}} \rangle$ starts the approach to its infinite temperature value $\langle H_{\text{static}} \rangle_{T=\infty} = 0$.

To highlight the prethermalization timescale $\tau_{D_{\text{eff}}}$, a different diagnostic is needed. In Fig. 3.1b, we use the second Rényi entropy of the leftmost three spins of the system $S_2 = -\log_2 \text{tr}[\rho_{\text{sub}}^2]$, where tr describes the trace operation, $\rho_{\text{sub}} = \text{tr}_{\text{site } i>3}[\rho]$, and ρ is the density matrix describing the system. Our choice of a small subsystem ensures that S_2 is not biased by DMT (see discussion below and [?]). While the system begins in a product state with $S_2 = 0$, its entropy quickly approaches a plateau, consistent with the Gibbs state of D_{eff} at a temperature which matches the energy density [?]; the system has approached the prethermal state. The timescale at which this occurs corresponds to $\tau_{D_{\text{eff}}}$ and, indeed, we observe $\tau_{D_{\text{eff}}} \sim 1/J_{\text{local}}$ independent of

the frequency ω of the drive. At later times $t > \tau^*$, parallel to $\langle H_{\text{static}} \rangle / L$, S_2 begins the approach to the infinite temperature value of $S_2^{T=\infty} = 3$ bits.

3.2 Benchmarking DMT

To ensure the reliability of DMT in the simulation of Floquet dynamics, we compare it Krylov subspace evolution [15, 89, 185] at small system sizes ($L = 20$). But before we compare DMT to Krylov, we verify that the Krylov evolution gives accurate results. This analysis not only gauges the applicability of DMT, but also leads to insights into the nature of the Floquet heating.

Time evolution with DMT proceeds via two repeating steps, each acting on a matrix product representation of the system's density operator. The two steps are a TEBD-like approximation of the time evolution unitary and a truncation via DMT. In the TEBD-like step, we Trotter decompose the time evolution operator into a series of gates which we then apply to the MPDO. Because each gate application increases the bond dimension of the MPDO, we must truncate it back to a fixed maximum bond dimension, which we call χ . During the truncation step, DMT separates χ into two contributions: χ^{preserve} and χ^{extra} . $\chi^{\text{preserve}} = 2^\ell$ is used to guarantee the preservation of all observables on ℓ contiguous sites during the truncation [?], and we call ℓ the *preservation diameter*. χ^{extra} is then used to preserve the remaining correlations with largest magnitude. We emphasize that, although truncation does not directly affect ℓ -sized operators, their dynamics is affected by the truncation of larger sized operators via the evolution of the system.

3.2.1 Benchmarking Krylov subspace dynamics

In this section we study the error of the Krylov subspace method by comparing to exact diagonalization (ED) at small system size. For each system size $L \in \{4, 6, 8, 10\}$, we consider a random initial product state of spins aligned in the \hat{z} direction. We then compute the evolution of the system under driving frequencies $\omega \in \{5, 6, 7, 8, 9, 10\}$.

In Fig. 3.3, we show the difference between the two numerical methods for the quantities of interest in our study. In particular we consider energy density H_{static}/L , second Rényi Entropy S_2 , onsite operators ($\sigma_i^{\{x,y,z\}}$) and two-site operators ($\sigma_i^x \sigma_{i+1}^x$ and $\sigma_i^z \sigma_{i+1}^z$). In the top row of plots of Fig. 3.3, we consider the maximum error observed in the first $600/J$ time units of the evolution. Errorbars correspond to the standard deviation of the maximum error observed over 6 different initial states. We note that the maximum error in this regime

is not substantially affected by the system size.

It is perhaps more enlightening to estimate the rate of error growth: for a given simulation parameters, Krylov subspace dynamics should induce a small constant error per timestep. This observation is borne out by the data for $t < 600/J$. In the bottom row of Fig. 3.3 we show the maximal rate of error growth, for some quantity O :

$$R = \max_{t \leq 600} \frac{|O_{\text{ED}}(t) - O_{\text{Krylov}}(t)|}{t}. \quad (3.14)$$

This provides an estimate of the error growth as a function of the simulation time. Since the rate of growth is $\leq 10^{-7}$, we believe Krylov subspace methods are suitable for benchmarking DMT, even at long times.




Figure 3.3: Difference of measured quantities between exact diagonalization and Krylov subspace methods over $t \leq 600/J$ physical time. In particular we consider energy density $\langle H_{\text{static}} \rangle / L$, second Rényi Entropy S_2 , single body operators ($\sigma_i^{\{x,y,z\}}$) and two body operators ($\sigma_i^x \sigma_{i+1}^x$ and $\sigma_i^z \sigma_{i+1}^z$), in their respective columns. In the top row we consider the largest difference observed within the elapsed time. In the lower row, we consider the largest rate of error growth, defined as Eq. (3.14).

3.2.2 DMT and Krylov

In Figs. 3.2 and 3.4, we compare the time evolution under DMT and Krylov and observe good quantitative agreement for $t > \tau_{D_{\text{eff}}}$, even at small bond dimension $\chi = 64$. In this regime, the system is close to a local thermal state and therefore has a natural low-bond-dimension MPDO representation. This behavior is clearest in the evolution of the energy density $\langle H_{\text{static}} \rangle / L$, as it is not sensitive to the prethermalization process (Fig. 3.2a). Moreover, the

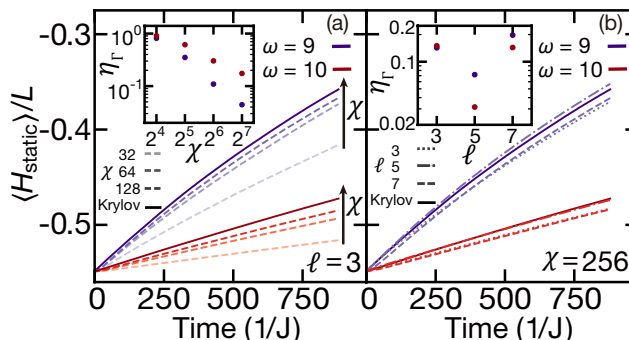


Figure 3.4: (a) Energy density at bond dimension $\chi = 32, 64, 128$ and (b) the preservation diameter $\ell = 3, 5, 7$. Insets: relative error in the heating rate $\eta_\Gamma = |\Gamma_{\text{DMT}} - \Gamma_{\text{Krylov}}| / \Gamma_{\text{Krylov}}$, in which Γ is defined by $\frac{d}{dt} \langle H_{\text{static}} \rangle = -\Gamma \langle H_{\text{static}} \rangle$.

same late-time behavior emerges when studying local observables, as well as the second Rényi entropy S_2 (Fig. 3.2b,c).

Closer inspection, however, shows that the small long-time disagreement increases with the frequency of the drive. This trend reflects the nature of Floquet heating. As the frequency increases, absorbing an energy quantum from the drive requires the rearrangement of a larger number of sites [6, 7, 155]. However, the truncation in DMT, limited by χ , often destroys the necessary longer ranged correlations ($\langle ab \rangle - \langle a \rangle \langle b \rangle \mapsto 0$ for operators a and b supported a distance $r > \ell$) suppressing the heating more severely at larger driving frequencies.

In Fig. 3.4, we leverage this very high-frequency regime to highlight the convergence properties of DMT with the bond dimension χ and the preservation diameter ℓ . As expected, increasing χ at fixed ℓ improves the accuracy of DMT since the amount of information preserved during each truncation is greater (Fig. 3.4a). Curiously, tuning ℓ at fixed χ can also affect the accuracy, despite not increasing the amount of preserved information (Fig. 3.4b). This suggests that, by carefully choosing which operators to preserve during truncation, one can achieve a high accuracy in DMT while keeping a low total bond dimension χ .

Having studied the late-time heating via the energy density, we now focus on the early-time dynamics via local observables and the second Rényi entropy. Unlike energy density, which is insensitive to the prethermalization dynamics, both local observables and S_2 can accrue significant errors for $t < \tau_{D_{\text{eff}}}$ at small bond dimension. This reflects the importance of long-range coherences in the early-time thermalization dynamics. A clear example is the dips in S_2 that

DMT fails to capture (arrows in Fig. 3.2c). We interpret these dips as the growth of operators bouncing back at the other edge of the system — DMT is unable to correctly capture the dynamics of the intermediate long-range operators and thus misses their return, resulting in a larger entropy value. We note that, by increasing χ , DMT better captures longer ranged operators and better captures the resonant dip, Fig. 3.2c inset. Curiously, although both S_2 and local observables are affected by long-range coherences, DMT appears to better capture the entropy dynamics. We believe this arises from the higher sensitivity of local observables to the details of the underlying interactions.

One might expect the large early-time deviations to lead to equally large late-time deviations. This is not what we observe. In particular, despite a considerable error in the early-time evolution under DMT, the system approaches the correct prethermal state at $\tau_{D_{\text{eff}}}$, and exhibits the correct energy density, local observables and entropy. To understand this behavior, we note that the prethermal Gibbs state can be fully determined by the system’s energy density, which is conserved until $\tau^* \gg \tau_{D_{\text{eff}}}$. Because both exact evolution and DMT conserve D_{eff} , the dynamics are constrained to the same subspace of the Hilbert space. As a result, although the precise agreement between the dynamics is limited by the finite bond dimension χ , as the evolution explores this subspace, both methods reach the same equilibrium state. For $t > \tau^*$, the system absorbs energy from the drive, and evolves between different thermal states of D_{eff} at the heating rate (for further evidence see [?]). In this regime, DMT accurately simulates observables to the extent it accurately simulates the heating process.

Finally, we explain how our choice to study S_2 on only three sites is motivated by the choice of preservation diameter $\ell \geq 3$. During the truncation procedure, by ensuring that any operator acting on ℓ contiguous sites is preserved, DMT also preserves the corresponding density matrix. As a result, truncation via DMT does not affect the entropy if the subsystem considered is smaller or equal to ℓ . As such, errors in the entropy arise from errors in longer ranged operators propagated via the system’s dynamics. This ensures that by measuring the accuracy of S_2 , we are measuring the accuracy of the dynamics rather than the impact of the truncation on the density matrix. In fact, when considering a subsystem larger than ℓ , the truncation biases the measured entropy. Each truncation step reduces the purity of the system, generating thermal-like entropy which leads to an extensive rate of entropy growth (for a detailed discussion see [?]).

3.2.3 Entropy in DMT

In the main text, we show the evolution of the entropy of the leftmost three sites. Here we provide more details on the entropy of subsystems with different sizes that motivates that choice for the main text.

3.2.4 Effect of Trotter step size on DMT numerics

In the main text, we considered the convergence of DMT with respect to the bond dimension χ and the size of preserved operators ℓ . Here we complement that analysis by considering the convergence in the size of the Trotter step. We quantify the error by measuring the average error

$$\delta_{\langle \hat{O}_i \rangle} \equiv \sqrt{L^{-1} \sum_i (\langle \hat{O}_i \rangle_{\text{DMT}} - \langle \hat{O}_i \rangle_{\text{Krylov}})^2}, \quad (3.15)$$

i.e., the error of a local observable averaged over all sites.

In Fig. 3.5, we take $\hat{O}_i \in \{\sigma_i^z \sigma_{i+1}^z, \sigma_i^x \sigma_{i+1}^x, \sigma_i^x\}$, the three local observables that contribute to energy H_{static} . By decreasing the Trotter step size from $4/J$ to $10/J$, we observe an improvement of results, especially during the late-time heating. However, the simulation does not benefit from further decreasing the step size from $10/J$ to $20/J$ at late time, although the error at very early-time decreases. This is because to apply truncations too frequently will destroy the many-body correlators more severely. We therefore use Trotter step $10/J$ throughout this work.

3.2.4.1 Page-like correction at late time

For a subsystem with size $L_{\text{sub}} \sim L/2$, the bipartite entanglement entropy approaches the maximal value of L_{sub} bits in DMT simulations, but a smaller value in Krylov subspace simulations (see Fig. 3.6). This discrepancy arises because the Krylov method guarantees that the system remains in a pure state, while in DMT the MPDO representation is entangled with a notional bath arising from the truncation procedure. The difference between the two corresponds to a Page-like correction. At $L_{\text{sub}} = L/2$ and infinite temperature, this correction is exactly $\log 2$ (1 bit), in agreement with the theoretical prediction [142]. As one decreases L_{sub} , the correction decreases exponentially, and the two methods agree better and better.




Figure 3.5: Convergence of local observables with respect to Trotter steps ($L = 20$, $\omega = 8$, $\chi = 128$, $\ell = 3$). (a) $\langle \sigma_i^z \sigma_{i+1}^z \rangle$, (b) $\langle \sigma_i^x \sigma_{i+1}^x \rangle$, and (c) $\langle \sigma_i^x \rangle$ are all the three terms that have non-zero contribution to energy density.

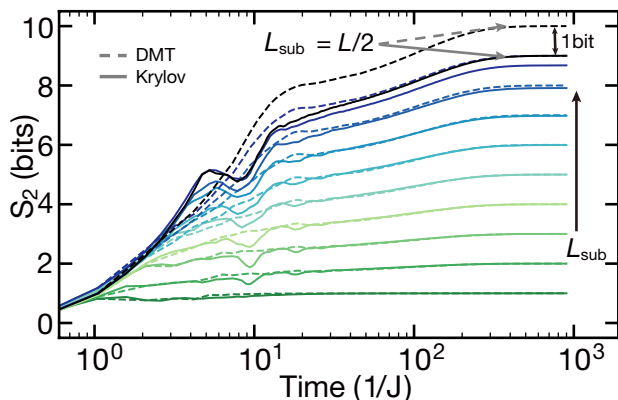


Figure 3.6: Late-time second Rényi entropy of subsystems with different sizes ($L = 20$, $\omega = 6$). The subsystem sizes L_{sub} for the curves from the bottom to top are $\{1, 2, 3, 4, 5, 6, 7, 8, 9, 10\}$.

3.2.4.2 Early-time behaviors

For $t < \tau_{D_{\text{eff}}}$, the system is well described by the time-independent interacting Hamiltonian D_{eff} . For an initial product state, the entropy of a subsystem is expected to increase linearly with time. Using Krylov subspace methods, we indeed observe the linear increase of entropy at early time. Curiously, the system often exhibits some non-monotonicity at intermediate times (from long-range coherences), but it eventually approaches its prethermal value at




Figure 3.7: Early-time second Rényi entropy of subsystems with different sizes ($\omega = 10$). (a) Small system with $L = 20$, and the subsystem sizes L_{sub} for the curves from the bottom to top are $\{1,2,3,4,5,6,7,8,9,10\}$. (b) Large system with $L = 100$, and the subsystem sizes L_{sub} for the curves from the bottom to top are $\{1,9,17,25,33,41,49\}$. The black straight line sets the upper bound of the entropy, of which the rate is extracted from Krylov data with $L = 20$.

$\tau_{D_{\text{eff}}}$ (Fig. 3.7a). We can extract the early-time entropy production rate Γ_{early}^S . Importantly, $t \cdot \Gamma_{\text{early}}^S$ sets an upper bound on the of entropy production (black line in Fig. 3.7).

However, DMT fails to capture both the effect of the many-body coherences, as well as the previously mentioned bound in entropy growth. On the one hand, the truncation destroys the long-range coherences necessary to capture the many-body revivals observed when using Krylov, resulting in a smoothed entropy curve. On the other hand, at subsystem sizes greater than the preservation diameter ℓ , the entropy in DMT can exceed the upper bound $t \cdot \Gamma_{\text{early}}^S$ (Fig. 3.7b). This is because the truncation in DMT can convert some entanglement entropy to thermal entropy, which does not care about the subsystem boundary. The entropy of a subsystem can then increase with the number of bonds truncated, so in larger subsystems we observe a higher rate of entropy growth (Fig. 3.7b).

3.3 Analysis and results

Having established that DMT accurately captures the late-time thermalization of quantum systems, we now apply it in the study of the late-time emergent

hydrodynamics in large spin chains ($L = 100$), beyond what can be achieved with previous methods.

Our ultimate goal is to write down a phenomenological description of the state and its dynamics, and to verify numerically that it describes our model.

The phenomenology is easily written down. In general, the dynamics of the local energy density follows the continuity equation: $\partial_t \epsilon(x, t) = \partial_x j(x, t) + \dot{Q}(x, t)$, where j is the energy current and \dot{Q} describes heating. Near equilibrium, the energy current is $j \propto \partial_x T$, where T is the local temperature. In contrast to homogeneous systems where $\partial_x T \propto \partial_x \epsilon$, the spatial inhomogeneity induced by the half-system drive forces the current to depend not only on $\partial_x \epsilon$, but also the position in the chain; thus a small correction in j is needed. At the same time, motivated by the exponential heating of the global drive, we expect the heating rate at the driven sites to be $\dot{Q} = -\epsilon(t)/\tau_{\text{local}}^*$, where τ_{local}^* is set by the global heating timescale τ_E^* .

We combine both these effects — heating and diffusion across the chain — in a simple hydrodynamical equation for the local energy density:

$$\partial_t \epsilon(x, t) = D(\epsilon) \partial_x^2 \left([1 + \eta g_1(x)] \epsilon(x, t) \right) - g_2(x) \frac{\epsilon(x, t)}{\tau_{\text{local}}^*}, \quad (3.16)$$

where $D(\epsilon)$ is the energy density dependent diffusion coefficient, while $g_1(x)$ and $g_2(x)$ characterize the regions where the drive affects diffusion and heating, respectively (the magnitude of the latter is controlled by a small parameter η). We set $g_1(x) = g_2(x) = \frac{1}{2} + \frac{1}{2} \tanh[(x - L/2)/\xi]$, a smoothed step function, and $\xi = 5$. (Our results are not sensitive to the particular choice of ξ , as long as $\xi \ll L$.)

Verifying this phenomenology is less straightforward. To do so, we consider two protocols: Fig. 3.8 offers a précis of the process. After verifying that the system is in fact well-described by a Gibbs state with time-varying temperature (Sec. 3.3.1, not shown in Fig. 3.8), we verify that the system displays an appropriate Floquet heating rate under both a half-system drive and a global drive (Fig. 3.8a).

We then turn to the details of the dynamics of the local energy density. To provide a baseline for comparison, we measure the diffusion coefficients of the static Hamiltonian H_{static} (Sec. 3.3.3, Fig. 3.8b). We then apply the half-driving protocol and fit the results to the modified diffusion equation (3.16)




Figure 3.8: (a) Heating rate τ^* extracted from energy density and subsystem entropy for $L = 20, 100$, as well as global and half-system drive. (b) Evolving with the static Hamiltonian H_{static} , an initial spatial variance in energy density will decay. The diffusion coefficient (at different energy densities) can be extracted from the exponential decay of the amplitude of Fourier modes. The amplitudes are normalized. (c) Dynamics of the energy density. By starting with the data at some initial state (e.g. $t = 200/J$) the hydrodynamical description Eq. (3.16) (dashed black curve) is able to correctly capture the energy density at late times.

(Fig. 3.8c, Sec. 3.3.4).

3.3.1 Approach to Gibbs Ensemble

To show that the system approaches a Gibbs ensemble at late time, we compare the expectation value of local observables of the late-time Floquet state and the Gibbs state of the static Hamiltonian H_{static} , which approximates the prethermal Hamiltonian D_{eff} to order ω^{-1} . By doing imaginary-time evolution with DMT, we obtain the Gibbs states of the H_{static} at different temperatures. The inset of Fig. 3.9 shows the averaged energy density $\langle H_{\text{static}} \rangle$ at different temperatures; similarly we can obtain other physical quantities as a function of temperature. This provides the map between energy density and the expectation value of other observables in the Gibbs state, allowing us to directly compare the late-time Floquet state to the Gibbs state at the same averaged energy density. As shown in Fig. 3.9, the two states have the same entropy and local observables, given the same energy density, indicating that the Floquet system can be described by a Gibbs ensemble with respect to the prethermal

Hamiltonian (for local observables).

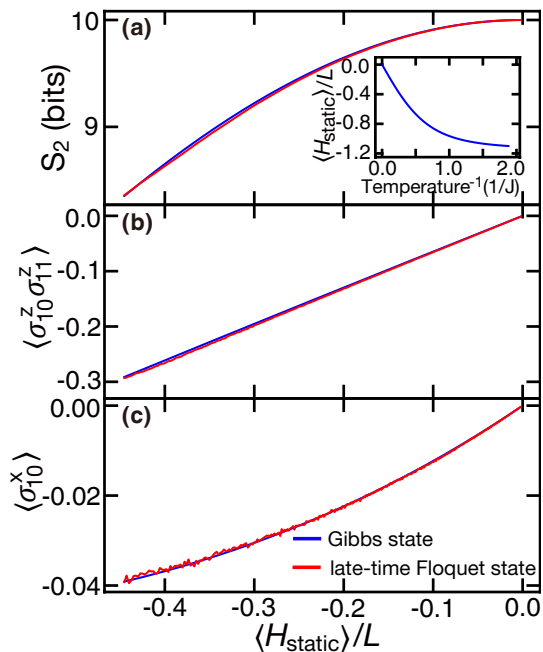


Figure 3.9: Physical quantities as a function averaged energy density for Gibbs states and late-time Floquet states ($L = 20$, $\omega = 6$). (a) second Rényi entropy of the half chain, (b) a two-site local observable, and (c) an one-site local observable. Inset: The averaged energy density as a function of inverse temperature.

How can these results be consistent with our claim in the main text that the relationship between effective temperature and (e.g.) the energy density of H_{static} is itself frequency dependent? Essentially, consistency requires that the particular local observables of Fig. 3.9 depend weakly on frequency. We find that they depend on frequency only at second order or higher in ω^{-1} , because those observables and H_{static} are even under global spin flip, while our drive is odd.

To see these symmetry considerations play out, suppose the system is in a Gibbs state

$$\rho(t; \omega) = \frac{1}{Z} e^{-\beta(t) D_{\text{eff}}(\omega^{-1})}, \quad (3.17)$$

where the effective temperature $T(t) = 1/\beta(t)$ is given by

$$\langle D_{\text{eff}}(t) \rangle \equiv Z^{-1} \text{tr}[D_{\text{eff}}(\omega^{-1}) e^{-\beta(t) D_{\text{eff}}(\omega^{-1})}] \quad (3.18)$$

with $Z = \text{tr}[e^{-\beta(t) D_{\text{eff}}(\omega^{-1})}]$. For compactness we drop the time dependence of

the temperature. Expand D_{eff} in powers of ω^{-1} :

$$D_{\text{eff}} = H_{\text{static}} + \omega^{-1}D' + O(\hbar^2\omega^{-2}). \quad (3.19)$$

Per Ref. [150]:

$$\begin{aligned} \omega^{-1}D' &= \frac{1}{T} \int_0^T dt i \int_0^t dt' [H_{\text{drive}}(t'), H_{\text{static}}] \\ &= \frac{1}{T} \int_0^T dt i \int_0^t dt' v(t) \sum_{jk} [h_z \sigma_j^z + h_y \sigma_j^y, J \sigma_k^z \sigma_{k+1}^z + J_x \sigma_k^x \sigma_{k+1}^x + h_x \sigma_k^x] \\ &= i \frac{\pi \omega^{-1}}{2} \sum_{jk} [h_z \sigma_j^z + h_y \sigma_j^y, J \sigma_k^z \sigma_{k+1}^z + J_x \sigma_k^x \sigma_{k+1}^x + h_x \sigma_k^x] \\ &= \pi \omega^{-1} \sum_k [-h_z J_x (\sigma_k^y \sigma_{k+1}^x + \sigma_k^x \sigma_{k+1}^y) - h_z h_x \sigma_k^y - h_y J (\sigma_k^x \sigma_{k+1}^z + \sigma_k^z \sigma_{k+1}^x) \\ &\quad + h_y J_x (\sigma_k^z \sigma_{k+1}^x + \sigma_k^x \sigma_{k+1}^z) + h_y h_x \sigma_k^z] . \end{aligned} \quad (3.20)$$

It is immediately apparent that D' is odd under a π rotation about the x axis, while H_{static} is even—more specifically, if

$$X = \prod_j \sigma_j^x \quad (3.21)$$

then

$$\begin{aligned} X H_{\text{static}} X &= H_{\text{static}} \\ X D' X &= -D' , \end{aligned} \quad (3.22)$$

so

$$0 = \text{tr}[D' H_{\text{static}}^n] . \quad (3.23)$$

Then the partition function is

$$\begin{aligned}
Z &= \text{tr} \left[e^{-\beta(H_{\text{static}} + \omega^{-1}D')} \right] + O(\omega^{-2}) \\
&\approx \text{tr} \left[e^{-\frac{\beta\omega^{-1}}{2}D'} e^{-\beta H_{\text{static}}} e^{-\frac{\beta\omega^{-1}}{2}D'} \right] + O(\omega^{-2}) \\
&\approx \text{tr} \left[\left(1 - \frac{\beta\omega^{-1}}{2}D' \right) e^{-\beta H_{\text{static}}} \left(1 - \frac{\beta\omega^{-1}}{2}D' \right) \right] + O(\omega^{-2}) \quad (3.24) \\
&\approx \text{tr} \left[e^{-\beta H_{\text{static}}} (1 - \beta\omega^{-1}D') \right] + O(\omega^{-2}) \\
&= \text{tr} \left[e^{-\beta H_{\text{static}}} \right] + O(\omega^{-2}) \\
&= Z_0 + O(\hbar^2\omega^{-2})
\end{aligned}$$

with $Z_0 = \text{tr} \left[e^{-\beta H_{\text{static}}} \right]$ the partition function of the static Hamiltonian.

Consider now some (local) operator O_j . With this fact about the partition function in mind its expectation value in the Gibbs state ρ is

$$\begin{aligned}
\text{tr} O_j \rho &= Z_0^{-1} \text{tr} \left[e^{-\frac{\beta\omega^{-1}}{2}D'} e^{-\beta H_{\text{static}}} e^{-\frac{\beta\omega^{-1}}{2}D'} \times O_j \right] + O(\omega^{-2}) \\
&= Z_0^{-1} \text{tr} \left[\left(1 - \frac{\beta\omega^{-1}}{2}D' \right) e^{-\beta H_{\text{static}}} \left(1 - \frac{\beta\omega^{-1}}{2}D' \right) \times O_j \right] + O(\omega^{-2}) \\
&= Z_0^{-1} \text{tr} \left[e^{-\beta H_{\text{static}}} O_j \right] - \frac{\beta\omega^{-1}}{2} Z_0^{-1} \text{tr} \left[e^{-\beta H_{\text{static}}} \{D', O_j\} \right] + O(\omega^{-2}), \quad (3.25)
\end{aligned}$$

where $\{\cdot, \cdot\}$ corresponds to the anti-commutator.

If O_j is even under X , as the operators of Fig. 3.9 are, $\text{tr} \left[e^{-\beta H_{\text{static}}} D' O_j \right] = 0$ and, to first order in ω^{-1} , O_j takes the same expectation value it has in the H_{static} Gibbs state:

$$\text{tr} \left[\rho O_j^{\text{even}} \right] = Z_0^{-1} \text{tr} \left[e^{-\beta H_{\text{static}}} O_j^{\text{even}} \right] + O(\omega^{-2}). \quad (3.26)$$

If O_j is odd under X , then

$$\text{tr} \left[\rho O_j^{\text{odd}} \right] = -Z_0^{-1} \frac{\beta\omega^{-1}}{2} \text{tr} \left[e^{-\beta H_{\text{static}}} \{D', O_j^{\text{odd}}\} \right] + O(\omega^{-2}). \quad (3.27)$$

Here we note that $Z_0^{-1} \text{tr} \left[e^{-\beta H_{\text{static}}} \{D', O_j^{\text{odd}}\} \right]$ corresponds to the sum of expectation values of one and two-body operators. Taking $O_j^{\text{odd}} = \sigma_k^y$, the significant terms in D' are $h_z J_x (\sigma_{k+1}^x \sigma_k^y + \sigma_k^y \sigma_{k+1}^x) + h_z h_x \sigma_k^y$. Using, $\beta \sim 0.2$, $h_z = 0.13$, $h_x = 0.21$, $J_x = 0.75$, $\omega = 6$ and the data from Fig. 3.9, we estimate




Figure 3.10: Comparison of the expectation value of σ^y for different sites on the half-driven chain, at frequency $\omega = 6$ and using the parameters of the main text. In red, the site is located within the driven side of the chain, where the effective Hamiltonian is modified by the drive. In blue, the site is located within the undriven side of the chain, where the effective Hamiltonian is only modified by the drive at very large order.

$\text{tr} \rho \sigma_k^y \sim \beta \pi \omega^{-1} (2h_z J_x \langle \sigma_k^x \rangle + h_z h_x) \sim 2.2 \times 10^{-3}$. In fact when we consider our half-driven chain and compare σ^y for a site in the undriven part ($D_{\text{eff}} = H_{\text{static}}$) to a site in the driven part ($D_{\text{eff}} = H_{\text{static}} + \omega^{-1} D' + \mathcal{O}(\omega^{-2})$), we find that in the driven part $\langle \sigma^y \rangle \sim 3 \times 10^{-3}$ whereas in the undriven region $\langle \sigma^y \rangle \sim 0$ (Fig. 3.10), in agreement with our estimate.

3.3.2 Heating

To ensure our results are not affected by finite-size effects, we revisit the global drive protocol in Figs. 3.1a,b. As in the $L = 20$ case shown in Figs. 3.2a,b, we observe a prethermal plateau in S_2 followed by the late-time heating of both $\langle H_{\text{static}} \rangle / L$ and S_2 toward the infinite temperature state. We characterize the late-time dynamics by extracting the heating timescale associated with their evolution. More specifically, given the observed exponential decay of both $\langle H_{\text{static}} \rangle / L$ and S_2 to their infinite-temperature values, i.e. $|\langle H_{\text{static}} \rangle / L| \propto e^{-t/\tau_E^*}$ and $(S_2^{T=\infty} - S_2) \propto e^{-2t/\tau_S^*}$, we extract τ_E^* and τ_S^* as a measure of the heating timescale τ^* [?]. Crucially, τ_E^* and τ_S^* agree with one another for both system sizes, Fig. 3.8a, demonstrating they capture the same heating process and are not affected by finite size effects. By repeating the same procedure for different frequencies of the drive, we observe the expected exponential dependence in frequency, $\tau_{E,S}^* \propto e^{\omega/J_{\text{local}}^{E,S}}$, suggesting previous bounds are tight [4, 6, 7, 57, 126, 155]. The extracted effective local energy scales, $J_{\text{local}}^E = 1.21 \pm 0.04$ and $J_{\text{local}}^S = 1.16 \pm 0.04$, are in good agreement with the microscopic onsite energy scale of H_{static} ($\|\hat{h}_i\| \simeq 1.26$, where $\hat{h}_i = J \sigma_i^z \sigma_{i+1}^z + J_x \sigma_i^x \sigma_{i+1}^x + h_x \frac{(\sigma_i^x + \sigma_{i+1}^x)}{2}$).

is the local Hamiltonian on each bond).

Having investigated Floquet heating with DMT, we now move to study the emergent diffusive behavior of local energy density. By performing imaginary-time evolution, we prepare the initial state to be a thermal state (Gibbs ensemble) with a small spatial variation in energy (taken to be different Fourier modes) [?]. As the system evolves and approaches the global equilibrium state, the amplitude of this spatial variation decays. Indeed, we observe the exponential decay of the amplitude of these modes. Crucially, this decay timescale grows as q^2 , where $2\pi/q$ is the wavelength of the Fourier mode, corroborating the diffusive nature of the dynamics [?]. By also changing the temperature of the initial Gibbs ensemble, the diffusion coefficient D can be studied as a function of average energy density of the system ϵ , Fig. 3.8b. We note our method, near infinite temperature ($\epsilon = 0$), matches independent calculations of the diffusion [162].

The ability to correctly simulate Floquet heating and diffusive behaviors in large-scale quantum systems enables us to study more complex late-time hydrodynamics. We now attempt to combine these two aspects by investigating the half-system driving protocol, where the drive only adds energy to the right side of the system. As a result, the heating timescale τ^* is twice of that of the global drive, Fig. 3.8a. In this setup, the dynamics of the local energy density is much richer than in the globally driven case: due to the inhomogenous heating, energy must be transported across the spin chain for the system to approach the infinite temperature state, Fig. 3.1c. As such, both heating and transport are crucial to capturing the late-time equilibration of the system. Moreover, by changing the frequency of the drive, one directly modifies the heating rate exponentially without much change to transport, thus controlling the relative importance between the two.

3.3.3 Extracting diffusion coefficients of a spatially uniform, static Hamiltonian

In this section we discuss the method used to extract the diffusion constant for a spatially uniform Hamiltonian H , which is a sum of local terms:

$$H = \sum_{\text{site } i=0}^{L-1} h_i . \quad (3.28)$$

We present the numerical experiment performed using DMT and Krylov subspace methods, and how the diffusion constant can be extracted from the obtained data.

3.3.3.1 Numerical Experiment

To properly probe diffusive behavior it is imperative that the system is perturbed around an equilibrium (i.e. thermal) state of H . More specifically, we want to initialize the system in a thermal state of $H + \eta H_{\text{perturb}}$, where the form of H_{perturb} controls the type of perturbation imposed, while η controls its strength.

Since we are interested in studying the diffusion of energy when evolved H , we want the perturbations to correspond to the eigenmodes of the diffusion equation, spatial oscillations of the energy density. We then consider a family of $H_{\text{perturb}}^{[k]}$ which generate the k -th mode:

$$H_{\text{perturb}}^{[k]} = \sum_{\text{site } i=0}^{L-1} h_i \cos \frac{ki\pi}{L-1}. \quad (3.29)$$

In DMT, the thermal state can be straightforwardly generated by performing imaginary time evolution on the infinite temperature state $\rho_{T=\infty} \propto I$:

$$\rho_\beta = Z^{-1} \exp \{ -\beta [H + \eta H_{\text{perturb}}] \}. \quad (3.30)$$

Because Krylov subspace methods can only treat pure states, it is impossible to directly compute expectations of the thermal state. Nevertheless, expectation values over the thermal density matrix ρ_β can be obtained by averaging over initial states, which are then imaginary time evolved:

$$\text{tr } O \rho_\beta = \frac{1}{D} \sum_{i=1}^D \langle \psi_i | \rho_\beta^{1/2} O \rho_\beta^{1/2} | \psi_i \rangle \approx \frac{1}{N_{\text{ave}}} \sum_{i=1}^{N_{\text{ave}}} \left[\langle \psi_i | \rho_\beta^{1/2} \right] O \left[\rho_\beta^{1/2} | \psi_i \rangle \right]. \quad (3.31)$$

Due to the large size D of the Hilbert space, we cannot perform the entire calculation. Instead we approximate it by averaging over N_{ave} number of *random* initial states $|\psi_i\rangle$:

$$|\psi_i\rangle \propto \sum_{i=1}^D c_i |i\rangle, \quad c_i \text{ normally distributed complex variables}. \quad (3.32)$$

Due to quantum typicality, such random states behave as infinite temperature states (for local operators) [181], and so the number of N_{ave} need not be very large (we use $N_{ave} = 50$).

Once the initial state is generated, the system is time evolved with H , and the local energy h_i is calculated as a function of time evolved. We observe that the initial spatial profile of the local energy quickly decays and the system becomes spatially uniform due to the diffusion of the energy density.

3.3.3.2 Extraction of the diffusion constant

Consider a system with some conserved quantity $S = \sum_j s_j$ such that s_j are local operators. Moreover let $\dot{s}_j = i[H, s_j]$ be also local (as is guaranteed in for a local Hamiltonian). We call s_j local conserved quantities. In our case, $S = H$ and $s_j = h_j$, the local energy.

When H is spatially uniform and at equilibrium $s_j(t)$ will be constant for all sites (up to edge effects). As a result, we can measure the distance from equilibrium by

$$\mathcal{P}(t) = \sqrt{\sum_j (\langle s_j(t) \rangle - \bar{s})^2} \quad (3.33)$$

where $\bar{s} \equiv S/L$ is independent of time. The decay of this quantity provides a proxy for the diffusion coefficient: if the system is diffusive with diffusion coefficient D , then the decay rate of this quantity is given by the decay rate of the slowest non-zero diffusive mode. For a generic initial state, this corresponds to

$$\mathcal{P}(t) \propto \exp[-\pi^2 t D / L^2] \quad (3.34)$$

for $t \gg L^2 / (4\pi^2 D)$ (the decay rate of the second-slowest mode). The diffusion coefficient is extracted by fitting this long-time behavior of $\mathcal{P}(t)$.

Alternatively, we can probe that slowest mode directly, by exciting a particular diffusive mode and measuring its magnitude. This is most straightforwardly implemented by preparing the lowest diffusive mode $k = 1$, $s_j(t = 0) \approx \cos(j\pi/L) + C$ as the initial state and measuring the amplitude of the corresponding Fourier mode:

$$s_{q=\pi/L}(t) = \frac{2}{L} \sum_{j=0}^{L-1} \cos \frac{j\pi}{L-1} s_j(t) . \quad (3.35)$$




Figure 3.11: (a) The evolution of the first Fourier mode under H_{static} . (b) Decay of Fourier modes at large system size. β is chosen such that the averaged energy density is set to be $\bar{\epsilon} = -0.1$. (c) The decay rate of Fourier modes depends quadratically on the wavevector. The system size $L = 100$.

In this, the decay of $s_{q=\pi/L}(t)$ will be $\propto e^{-\pi^2 t D/L^2}$, from where D can be extracted. We note that the profile of $s_j(t)$ can also be fitted, with a least-square method, to the lowest Fourier mode. Both methods yield the same results.

Fig. 3.11a,b illustrate both methods, investigating $\mathcal{P}(t)$ and $s_{q=\pi/L}(t)$ for $H = H_{\text{static}}$. We see both the slowest mode and the sum over all modes decay exponentially with time; they have the same decay rate. Here we note that, at early times, we observe non-exponential behavior in the decay of $s_{q=\pi/L}(t)$ until a timescale $\sim 1/h_x$, the integrability breaking term of our system. It is at this timescale that we expect the interactions to induce the appropriate diffusive behavior.

Moreover, we can study the decay of higher Fourier modes by using the same methodology. In Fig. 3.11(c), we observe the quadratic dependence of the decay rate on the wavevector, supporting that the dynamics of local energy density is diffusive in our system.

3.3.3.3 Accuracy of extracted diffusion coefficients

The work of Kloss, Bar Lev and Reichman [119] and ongoing (unpublished) work of Leviatan et al. find that TDVP shows “false convergence”: it can converge very quickly in bond dimension—but to dynamics with an unphysical diffusion coefficient. We must therefore check that DMT shows the *correct* diffusive dynamics. In this section we compare the diffusion constant extracted between DMT and Krylov-space dynamics.




Figure 3.12: Decay of Fourier modes at small system size ($L = 20$) for (a) our model H_{static} near integrability, with the parameters used in the bulk of the paper, and for (b) our model far from integrability (b). In each The averaged energy density is set to be $\bar{\epsilon} = -0.1$.

Consider our static Hamiltonian

$$H_{\text{static}} = \sum_{i=1}^{L-1} [J\sigma_i^z\sigma_{i+1}^z + J_x\sigma_i^x\sigma_{i+1}^x] + h_x \sum_{i=1}^L \sigma_i^x. \quad (3.36)$$

In the main text we used $\{J, J_x, h_x, h_y, h_z\} = \{1, 0.75, 0.21, 0.17, 0.13\}$, similar to previous work [150]. Because we are considering nearest neighbor interactions, the only integrability breaking term is h_x , leading to a naïve estimate for the scattering length of $\lambda \sim J/h_x \simeq 5$. As a result, observing diffusion at small system sizes is difficult. Fig. 3.12 (left) highlights this difficulty. Nevertheless, we observe good agreement between DMT and Krylov in the dynamics. We believe that DMT artificially increases a dephasing rate for the model’s quasi-particles; this explains the gradually increasing discrepancy between the DMT and Krylov simulations.

To check how well DMT can capture diffusion, we increase $h_x = 1.03$, decreasing the scattering lengthscale and making small system sizes more amenable to studies of diffusion. Indeed, Fig. 3.12 clearly demonstrates the agreement between the two methods, and as a result, the ability of DMT to probe the diffusive physics.

3.3.4 Classical diffusion equation

With the previously extracted parameters τ_E^* and $D(\epsilon)$, we can calculate the hydrodynamical evolution by Eq. 3.16, and compare it with the DMT simulation of the microscopic model. We leave η to be the only fitting parameter, constant across the entire evolution, but dependent on the frequency of the drive. Starting with an initial product state, we let the system reach local equilibrium by evolving it to time $t_0 = 200/J$. Choosing the resulting energy

density profile as the initial state of Eq. 3.16, we observe excellent agreement of the remainder of the evolution, Fig. 3.8c. Using the same methodology, we reproduce the dynamics for different driving frequencies, corroborating our hydrodynamical description for the local energy density dynamics under the Floquet quantum model, Fig. 3.1c. Finally, we note that the fitting parameter η shows the expected negative dependence on the driving frequency ω ; at higher frequency, D_{eff} on the right side of the chain approaches H_{static} and the spatial inhomogeneity in D_{eff} is reduced.

3.3.4.1 Derivation and approximation

In general, the equations governing the heat transport in classical systems are

$$\partial_t \epsilon(x, t) = \partial_x j(x, t) + q(x, t) \quad (3.37)$$

$$j(x, t) \propto \partial_x T(x, t). \quad (3.38)$$

The first equation is energy conservation, and the second equation reflects that a non-uniform temperature $T(x, t)$ will lead to a heat current $j(x, t)$.

By driving the right half of the chain, the conversion from energy density $\epsilon(x, t)$ to temperature $T(x, t)$, as well as the energy absorption $q(x, t)$ varies explicitly in position along the chain. In particular, the lowest order $\epsilon(x, t)$ correction to $T(x, t)$ yields a heat current $j(x, t) \propto \partial_x T(x, t) \propto \partial[(1 + \eta g(x))\epsilon(x, t)]$, in which the small parameter η characterizes the correction of the conversion from ϵ to T on the driven half and $g(x)$ captures the spatial profile of the inhomogeneity. In our model, $g(x)$ should be close to the step function $\Theta(x - L/2)$. The diffusion equation then reads

$$\partial_t \epsilon = D \partial_x^2 [(1 + \eta g(x))\epsilon] - q(x, t) \approx D \partial_x^2 [(1 + \eta g(x))\epsilon] - \frac{\epsilon}{\tau^*} g(x), \quad (3.39)$$

where we have used the fact that $q(x, t) \approx -g(x)\epsilon/\tau^*$ (for the driven part).

We remark that a spatially varying D_{eff} may also lead to another two modifications in the heat equation: a spatially dependent definition of ϵ , and spatially dependent diffusion constant D . A heat equation including all these modifications can be written as

$$\partial_t [(1 + \lambda g(x))\epsilon] = \partial_x \{ [D + \delta D g(x)] \partial_x [(1 + \eta g(x))\epsilon] \} - \frac{\epsilon}{\tau^*} g(x), \quad (3.40)$$

where λ , δD and η are all small numbers. We find that the dynamics in our

experiment is not sensitive to the inclusion of these terms, as they do not lead to qualitatively different terms in the equation.

In fact, we find that the η term is the most meaningful for well describing the dynamics than λ and δD . At early times, before the heating occurs ($\tau_{D_{\text{eff}}} < t < \tau^*$), we already observe a non-homogenous spatial profile of energy density, due to the temperature inhomogeneity induced by η . Moreover, since the spatial profile $g(x)$ is close to the step function $\Theta(x - L/2)$, a higher derivative of it will contribute to larger correction. Therefore, the η term, which has a second derivative, should be the most important correction.

3.3.4.2 Solving the Heat Equation

Since no energy can flow out of the boundary of the system, we require $0 = j(t)|_{x=0,L} \approx \partial_x[(1 + \eta g(x))\epsilon]|_{x=0,L}$. Considering that $g(x)$ remains constant deeply inside the driven and undriven parts, we can simplify the boundary condition as $\partial_x \epsilon|_{x=0,L} = 0$. This can be immediately achieved by considering the cosine series of the problem

$$\epsilon(x, t) = \sum_{n=0}^{\infty} f_n(t) \cos \frac{n\pi x}{L}. \quad (3.41)$$

The differential equation then becomes:

$$\begin{aligned} \sum_{n=0}^{\infty} \cos \frac{n\pi x}{L} \partial_t f_n(t) &= -D \sum_{n=0}^{\infty} \left[\frac{n\pi}{L} \right]^2 \cos \frac{n\pi x}{L} f_n(t) \\ &+ D\eta \sum_{n=0}^{\infty} \partial_x^2 \left[g(x) \cos \frac{n\pi x}{L} \right] f_n(t) \\ &- \frac{1}{\tau^*} \sum_{n=0}^{\infty} g(x) \cos \frac{n\pi x}{L} f_n(t). \end{aligned} \quad (3.42)$$

Integrating both sides with against $\cos \frac{k\pi x}{L}$ for $k \in \mathbb{N}$ yields (the second term on the right-hand side is integrated by parts):

$$\begin{aligned} \partial_t f_k(t) = & -f_k(t) D \left[\frac{k\pi}{L} \right]^2 \\ & - \frac{2}{L(1 + \delta_{k0})} \sum_{n=0}^{\infty} f_n(t) \left\{ D\eta \left[\frac{k\pi}{L} \right]^2 + \frac{1}{\tau^*} \right\} \\ & \times \int_0^L dx g(x) \cos \frac{k\pi x}{L} \cos \frac{n\pi x}{L} . \end{aligned} \quad (3.43)$$

The resulting equations can be cast in a vectorial form as

$$\partial_t \vec{f}(t) = M \vec{f}(t) , \quad (3.44)$$

where \vec{f} is the vector of the Fourier components and M describes the coupling between the modes in the right-hand side of Eq. (3.43). In practice, the magnitude of the Fourier modes decays very quickly with n , so we can consider only the first $N = 40$ and not incur significant error.

The time evolution of \vec{f} then becomes exactly

$$\vec{f}(t) = \exp \{ (t - t_0) M \} \vec{f}(t_0) . \quad (3.45)$$

3.3.4.3 Dynamics of the energy density

We now describe the procedure by which we can obtain the dynamics of the energy density at late times in the $L = 100$ system.

The hydrodynamical description holds only for systems near a local equilibrium. As such to ensure we have a correct initial state, we choose some initial time t_0 and Fourier transform the energy density profile at that time, thus obtaining $f_n(t_0)$.

We can now consider the evolution iteratively, i.e. to consider a series of times $\{t_n\}$ and evolve our state using Eq. (3.45) from t_{n-1} to t_n . Here, the heating timescale τ^* and the diffusion coefficient $D(\epsilon)$ are from the global heating rate and the diffusion under static Hamiltonian, respectively. For the energy-independent parameter η , we optimize its value to minimize the discrepancy between the two evolutions (DMT and hydrodynamical model). In particular, we characterize the discrepancy by the standard deviation averaged over all time slices, namely, $\sum_{t_n} \sqrt{\sum_x |\epsilon(x, t_n)_{DMT} - \epsilon(x, t_n)_{Hydro}|^2 / L}$.




Figure 3.13: Time evolution of the energy density profile starting with the lowest energy density state as the initial state. We observe great agreement with the DMT results. The extracted values of η are shown in Fig. 3.14.

In Fig. 3.13, we compare the energy density profile from the simulation with that arising from Eq. (3.45). Using the lowest energy density state as the initial state, we can then apply the above procedure to obtain the energy density dynamics for later times across a large frequency range. We also observe that η has a negative dependence on frequency ω as expected (Fig. 3.14), since its value is determined by the higher order corrections to D_{eff} in ω^{-1} , which goes smaller when the driving frequency increases. Moreover, by increasing bond dimension χ in DMT, we check the convergence of the energy density (Fig. 3.15).




Figure 3.14: Extracted η decreases as the driving frequency increases.

3.3.5 Hydrodynamics in large spin chains

3.4 Discussion

In this chapter, we used DMT to study the emergence of hydrodynamics in the late-time evolution of large-scale Floquet systems. We began by demonstrating the validity and the range of applicability of this new method by comparing it against Krylov subspace methods, showing excellent agreement of the late-




Figure 3.15: Convergence of energy density. (a) Direct comparison between different bond dimension χ . (b) The relative error of local energy density $\delta\epsilon = \epsilon_\chi - \epsilon_{\chi=180}$. We chose $\omega = 6$ and averaged energy density $\bar{\epsilon} = -0.2$ as an illustration. For other choice of parameters, we observe similar trends.

time thermalization dynamics. By working at large system size ($L = 100$), we independently characterize the heating and the diffusive properties of the dynamics of local energy density in our model. Using the half-system drive, we combine these two effects and describe the resulting late-time dynamics with a simple hydrodynamical model, Eq. (3.16).

HYDRODYNAMICS IN A DISORDERED LOW-FREQUENCY FLOQUET SYSTEM

In Ch. 3, we considered the dynamics of a clean Floquet system. Since the work of Deutsch [54] and Srednicki [200], we have believed generic isolated quantum systems to satisfy the eigenstate thermalization hypothesis. We expect that for times greater than some local thermalization time, systems satisfying that hypothesis will be well-described by a hydrodynamical model. That is, that local observables will be given by a Gibbs state with spatially varying Lagrange multipliers for the model’s conserved density (temperature, chemical potential, etc.). Those Lagrange multipliers in turn evolve by a partial differential equation, generally (at lowest order) a diffusion equation. The key result of Ch. 3 was that one could extend this phenomenological description to include Floquet heating—as long as one was careful to draw the diffusion coefficient from the right (effective prethermal) Hamiltonian.

But Basko, Aleiner, and Altshuler [19] (cf [159]) introduced a broad class of disordered systems—so-called *many body localized* systems—that do not satisfy the eigenstate thermalization hypothesis, and are not well-described by a diffusive hydrodynamical theory with any small number of local conserved quantities. Instead, these MBL systems have local conserved quantities: their Hamiltonians commute with local operators (l-bits or local integrals of motion [92, 193]). Consequently, the system retains a local memory of its initial state even at long times.

Many disordered models display a transition: they satisfy the eigenstate thermalization at small disorder, but are localized at large disorder. The fruit-fly example is the random-field Heisenberg model [146, 160], but other examples—including the the second-neighbor hopping model of [159], a variant of which we study here—are also known to display both phases. A number of phenomenological renormalization group schemes exist to describe this transition [72, 172, 220]; all predict that on the thermalizing side of the transition, the system will display subdiffusive hydrodynamics due to rare insulating regions.

But it is still not clear how far the Griffiths region extends. Do systems with MBL transitions display subdiffusion all the way to the clean limit? [108, 147]

find that it does; [144, 189, 246] find that it does not. We find one such model that does not; in fact, it displays a large diffusive regime. We take a robustly non-integrable model (a Floquet variant of the second-neighbor hopping model studied by Huse and Oganesyan [159]) and use DMT to probe its hydrodynamics. We find that this model displays a large diffusive region. In this region, we measure the diffusion constant and a characteristic ultraviolet length scale for the model, which we identify as a coherence length; curiously, this coherence length does not change with disorder.

4.1 Model

We use a Floquet variant of the model of [159]. When written with fermion operators, that model is

$$H = \sum_i \left[w_i n_i + V(n_i - 1/2)(n_{i+1} - 1/2) + (c_i^\dagger c_{i+1} + h.c.) + (c_i^\dagger c_{i+2} + h.c.) \right]. \quad (4.1)$$

We, like Huse and Oganesyan, take $V = 2$. Upon Jordan-Wigner transformation the model becomes

$$H = \sum_i \left[2(\vec{S} \cdot \vec{S} + [S_i^+ S_{i+1}^z S_{i+2}^- + h.c.]) + w_i S_i^z \right]. \quad (4.2)$$

In each case the disorder fields w_i are normally distributed about 0 with width W . The continuous-time model satisfies the ETH at $W = 0$, while Huse and Oganesyan find an ETH-MBL transition around $W = 6$.

The model (4.2) has two local conserved quantities, namely spin (particle number) and energy density. The coupling between the two conserved quantities leads to important nonlinear corrections to the hydrodynamics [156, 202]. We avoid this complication by breaking energy conservation: we time evolve by a series of unitaries

$$U(t) = (U_{\text{even}} U_{\text{odd}})^t, \quad t \in \mathbf{Z} \quad (4.3)$$

and stroboscopically sample. Spatially random Floquet models have been used as generic toy models for quantum thermalization[33, 111, 219] for similar reasons: they provide minimally-structured models.

The precise form we choose for $U_{\text{even,odd}}$ is motivated by computational convenience: essentially, we group nearest-neighbor sites and then perform a first-order even-odd Trotter decomposition of the continuous-time evolution

$U(t) = e^{-iHt}$. More precisely,

$$\begin{aligned}
 U_{\text{even}} &= \prod_{j=1}^{\lfloor L/4 \rfloor} \exp[-i \delta t h_{2j}] \\
 U_{\text{odd}} &= \prod_{j=1}^{\lfloor L/4 \rfloor} \exp[-i \delta t h_{2j+1}] ,
 \end{aligned}
 \tag{4.4}$$

where L is the system size and

$$\begin{aligned}
 h_j^{(1)} &= \frac{1}{2} \vec{S}_{2j-1} \cdot \vec{S}_{2j} + \frac{1}{2} [W_{2j-1} S_{2j-1}^z + W_{2j} S_{2j}^z] \\
 h_j^{(2)} &= \frac{1}{2} \vec{S}_{2j} \cdot \vec{S}_{2j} + [S_{2j-1}^+ S_{2j}^z S_{2j+1}^- + \text{h.c.}] \\
 &\quad + [S_{2j}^+ S_{2j+1}^z S_{2j+2}^- + \text{h.c.}] \\
 h_j &= \alpha_{2j-1} h_j^{(1)} + \alpha_{2j} h_{j+1}^{(1)} + h_j^{(2)}
 \end{aligned}
 \tag{4.5}$$

with α_l chosen to take care of the left and right boundaries:

$$\alpha_l = \begin{cases} \frac{1}{2} & l = 1 \text{ or } l = L - 1 \\ 1 & \text{else .} \end{cases}
 \tag{4.6}$$

δt is a parameter. We wish to choose δt large enough that a Magnus expansion far from convergence. If we choose δt too small then we may see an approximately local conserved Floquet energy density, which could lead to the nonlinear hydrodynamics we wish to avoid.

The Floquet model has an MBL transition around disorder width $W = 20$ (see Fig. 4.1).

4.2 Method

We use DMT, as described in Ch. 2, to simulate the time evolution of a random mixed initial state by the Floquet unitaries (4.3).

We take an initial state

$$\rho = \frac{1}{Z} \exp \left[\sum_j \mu_j S_j^z \right]
 \tag{4.7}$$

with μ_j random in $[-1, 1]$. We choose this initial state because it is not far from the identity—which makes us more confident of DMT’s efficacy—while containing variability at all wavelengths.




Figure 4.1: Gap ratio for the Floquet model of Sec. 4.1 as a function of disorder width at half-filling. Error bars show standard deviation of per-realization average gap ratio divided by square root of number of realizations.

Since the model in question (4.2) is three-local, we group each adjacent pair of physical sites (local onsite dimension 2) into a single effective site (local onsite dimension 4). The method then guarantees that no truncation will change any operator on six contiguous physical sites. This is a strong guarantee. In particular, truncation does not change

$$\langle S_j^z \rangle, \left\langle \frac{d}{dt} S_j^z \right\rangle, \dots, \left\langle \frac{d^5}{dt^5} S_j^z \right\rangle. \quad (4.8)$$

Consider, then, DMT at some small, fixed bond dimension. Thanks to this guarantee, we expect that DMT should approximate the model’s dynamics well in the hydrodynamical regime, where higher derivatives are very small—even if it requires much higher bond dimension to faithfully approximate short-to intermediate-time dynamics.

4.2.1 Convergence testing

One can crudely estimate truncation error by simulating the same system at two different bond dimensions (say χ_0 and $2\chi_0$) and comparing the expectation values of the observables at different bond dimensions. But as the $S_j^z(t)$ approach their (small) thermal values, this becomes less useful: we want to normalize the discrepancy by the site-to-site variation in one $S_j^z(t)$ or the other. (It is not sufficient to normalize by a *typical* value of $S_j^z(t)$, as this will just be given by the thermal value. Since we seek to measure deviations from this thermal value, we desire that our error be small compared to those deviations. Nor can we normalize by *individual* $S_j^z(t)$: as these cross zero, they will produce spurious divergences in the error.) We therefore define the **relative**

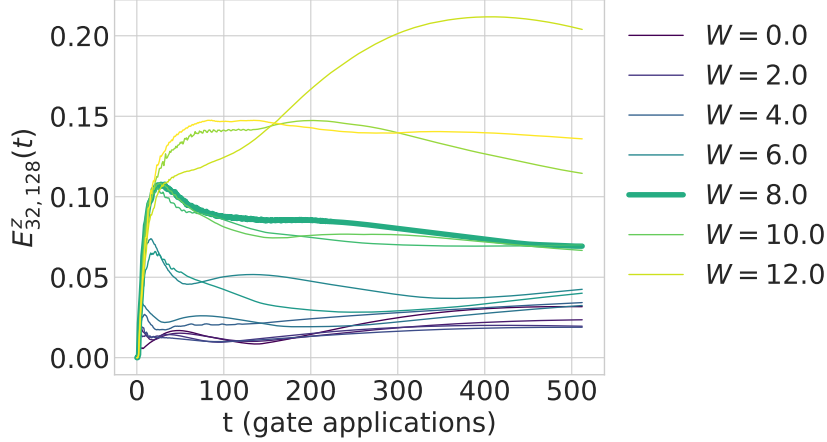


Figure 4.2: Relative convergence error (4.11) at a variety of disorder strengths, for system size $L = 64$. (We label only even disorder strengths, for compactness.) The thick line indicates $W = 8.0$, the largest disorder strength for which we judge the simulation to be converged.

convergence error between bond two bond dimensions χ_1, χ_2

$$E_{\chi_1, \chi_2}[S_j^z(t)] \equiv \frac{S_j^z(t, \chi = \chi_1) - S_j^z(t, \chi = \chi_2)}{\text{std}[S_j^z(t, \chi = \chi_2)]}. \quad (4.9)$$

where we normalize by the standard deviation

$$\begin{aligned} & \text{std}[S_j^z(t, \chi = \chi_2)] \\ & \equiv \left[\frac{1}{L} \sum_j S_j^z(t, \chi = \chi_2)^2 \right. \\ & \quad \left. - \frac{1}{L} \left(\sum_j S_j^z(t, \chi = \chi_2) \right)^2 \right]^{1/2}, \end{aligned} \quad (4.10)$$

which is a proxy for deviation from the thermal value and one quantity of physical interest. (We will consider the physics of this quantity in detail below.) In order to have a single measure, we take the RMS:

$$E_{\chi_1, \chi_2}^z(t) \equiv \left(\sum E_{\chi_1, \chi_2}[S_j^z(t)]^2 \right)^{1/2}. \quad (4.11)$$

We also call this the relative convergence error.

The error increases at a rate roughly independent of disorder strength W until a peak number of gate applications t_{worst} , after which the error decreases. t_{worst} . It is tempting to interpret t_{worst} as a local-thermalization time and the decay

of error after t_{worst} as confirmation of our intuition that our method captures hydrodynamics accurately (even if it fails to capture short- or intermediate-time dynamics). We leave to future work the careful diagnostics required to be confident in this interpretation.

We see that for $W \leq 8.0$, $E_{32,128}^z(t) \lesssim 0.1$; we judge the method to be trustworthy for these disorder strengths.

At first sight these results seem to contradict the previous chapter. There we compared DMT simulations of a Floquet system at high frequencies to quasiexact Krylov space simulations and found that DMT agrees very well with the Krylov space results at long times, but fails at short times ($t \sim 10/J$, J a coupling). We, on the other hand, find that DMT converges well (at least for clean systems) at all times.

The apparent discrepancy is due to the fact that the model of Ch. 3 is close to free-fermion integrability. Truncation with DMT is liable to change the nearly conserved free fermion occupation numbers; since these occupation numbers are important to the system's short-time dynamics (e.g. the revivals clearly visible in Fig. 3.2). This chapter's model, on the other hand, is far from any kind of integrability.

4.2.2 Comparison with exact simulations

But in light of the work of Kloss et al. [119], it is not enough to check convergence. They applied TDVP as in [133] and found that it could display false convergence: that as one increases the bond dimension, one sometimes sees the TDVP simulations converge quickly—but to a value that disagrees with exact results.

In this section we consider DMT on small systems ($L \leq 28$). Because the system size is small, we can compare DMT to exact simulations using MPSs of large bond dimensions (cf. Sec. ?? We consider two initial states: a product of random S^z eigenstates (chosen so the system is at half-filling) and a product state of rotated spins:

$$|\psi(0)\rangle = \bigotimes_{j=0}^{L-1} [\cos(\pi j/L) |\uparrow\rangle + \sin(\pi j/L) |\downarrow\rangle]. \quad (4.12)$$

(The latter is chosen to mimic the slowest mode of the continuous diffusion equation. Additionally, because it is not an eigenstate of $\sum_j S_j^z$, its dynamics

should better mimic the large-system limit.) In each case we compare DMT to MPS simulations with bond dimension $2^{L/2}$ —large enough that the MPS is never truncated (except to remove values below a numerical noise floor).

Fig. 4.3 shows the hydrodynamics extracted from product of random S^z eigenstates, while Fig. 4.4 shows that extracted from the product of superpositions (4.12). At first sight, Fig. 4.3 appears to show the false convergence that has bedeviled efforts to use TDVP to extract diffusion coefficients. Bond dimensions $\chi = 64$ and $\chi = 128$ give very nearly the same value—but this value is not the correct (exact) value.

But this difference between diffusion coefficients extracted from DMT and those extracted from exact dynamics is of the same order as the difference between diffusion coefficients extracted from the exact dynamics at different system sizes—which suggests that the physics DMT is failing to capture more than finite size effects. Turning to the second initial state, Eq. (4.12) and Fig. 4.4, we see more reasonable convergence properties— $\chi = 64, 128$ differ by an amount on the order of 1%—and still different “exact” values. We believe that the discrepancies are a combination of finite-size effects and imprecisions in our protocol.




Figure 4.3: Hydrodynamics starting from an S^z product state randomly chosen at half-filling. Left: decay of the distance from equilibrium (3.33) for $L = 26$. Center: diffusion coefficient, as calculated exactly (with large bond dimension matrix product states) and with DMT at three bond dimensions. Right: error in the DMT diffusion coefficients.




Figure 4.4: Hydrodynamics starting from the initial state (4.12). Left: decay of the distance from equilibrium (3.33) for $L = 26$. Center: diffusion coefficient, as calculated exactly (with large bond dimension matrix product states) and with DMT at three bond dimensions. Right: error in the DMT diffusion coefficients.

4.2.3 Quantity of interest

We directly measure $s_j(t) \equiv \langle S_j^z(t) \rangle$, but the time traces of onsite spins are not enlightening. Instead, we define a distance from equilibrium

$$\mathcal{P} \equiv \sum_j (s_j - \bar{s})^2, \quad (4.13)$$

where \bar{s} is the average spin

$$\bar{s} = \frac{1}{L} \sum_j s_j(t) = \frac{1}{L} \sum_j s_j(0). \quad (4.14)$$

Writing the s_j as a vector s , that distance from equilibrium is

$$\mathcal{P} = (s - \bar{s})^T (s - \bar{s}). \quad (4.15)$$

\mathcal{P} is the total weight on all of the system's decay modes: if $u_{(\alpha)}$ is a mode with decay rate $\gamma_{(\alpha)}$, i.e.

$$Au_{(\alpha)} = \gamma_{(\alpha)}u_{(\alpha)} \neq 0, \quad (4.16)$$

then

$$\mathcal{P} \equiv \sum_{\alpha} (s^T u_{(\alpha)})^2. \quad (4.17)$$

4.2.4 Disorder averaging

One can legitimately disorder-average any measurable quantity. The disorder average then characterizes the distribution of that quantity. This is not the only reasonable goal, however: one may wish to use the disorder average not to characterize the distribution at a given length, but to mimic a longer systems than one simulates, so the limit of many disorder realizations becomes the thermodynamic limit. Using a disorder average in this way, to probe large system sizes, requires some kind of phenomenological renormalization group scheme. One must disorder average in a way that mimics this scheme.

In our case, if we imagine connecting two disorder realizations with uniform resistances ρ_1, ρ_2 , the resulting segment has resistivity $\rho_{12} = \frac{1}{2}(\rho_1 + \rho_2)$. Since the diffusion coefficient is $D \propto \rho^{-1}$, this means that we should average diffusion coefficients by $D_{12} = [D_1^{-1} + D_2^{-1}]^{-1}$.

Consider now the distance from equilibrium \mathcal{P} in two regimes, divided by the Thouless time $t_{\text{Th}} \sim L^2/D$:

1. $1/J \ll t \ll t_{\text{Th}} \sim L^2/D$, when $\mathcal{P} \sim \frac{1}{\sqrt{Dt}}$. In this case $\mathcal{P} \propto D^{-1/2} \propto \sqrt{\rho}$, so we should disorder average by

$$\bar{\mathcal{P}} = \langle (P^{(r)})^2 \rangle^{1/2}. \quad (4.18)$$

2. $t_{\text{Th}} \ll t$, when $\mathcal{P} \sim e^{-(\pi^2 D/2L^2)t}$. In this case $P \sim e^{-t/\rho}$, so we should disorder average

$$\mathcal{P} = \exp \left[\langle (\ln P^{(r)})^{-1} \rangle^{-1} \right]. \quad (4.19)$$

In each case we denote by angle brackets a simple disorder average $\langle A \rangle = \frac{1}{N_{\text{reals}}} \sum_{r=1}^{N_{\text{reals}}} A^{(r)}$. Most of our results derive from the longest time behavior—that is, the decay of the slowest hydrodynamic mode. We therefore use the disorder averaging protocol 2.




Figure 4.5: Scaling collapse for the distance from equilibrium $\mathcal{P}(t)$ for 30 disorder realizations at $L = 64, W = 3.0$. Left: $\mathcal{P}^{(r)}$ for each of the 30 realizations. Middle, right: disorder average per the protocol 2. The kink in the $L = 512$ curve is an artifact of that disorder averaging protocol. The dashed black line shows $\mathcal{P}_1 \propto \frac{1}{\sqrt{t}}$. Error bars show 5th and 95th percentiles in a non-parametric bootstrap scheme; they do not include the simulation error $\mathcal{E}_{32,128} \lesssim 10\%$.

4.3 Results

We plot the distance from equilibrium $\mathcal{P}(t)$ for disorder width $W = 3.0$ in Fig. 4.5. There we see not only intermediate-time $\mathcal{P}_4 \sim t^{-0.5}$ scaling, but also a convincing finite-size scaling collapse $\tilde{t} = t/L^2$, which is strongly indicative of diffusive behavior. This collapse captures not only the long-time limit but also the crossover from the intermediate-time power law.

The scaling collapse does not work for $L = 16$. This is unsurprising. The scaling collapse should only work for system sizes $L \gg l_{UV}$, where l_{UV} is the UV cutoff of the diffusive hydrodynamics; because our Trotter gates are supported on four sites, one would expect the diffusive, hydrodynamical model to have a UV cutoff $l_{UV} \gtrsim 4$. In particular, the outermost four sites (two on the left and two on the right) are only affected by every other layer of gates; when $L = 16$, this is a significant portion of the system.

We can extract effective whole-system diffusion coefficients by considering the long-time behavior of the $\mathcal{P}^{(r)}$. In the long-time limit

$$\lim_{t \rightarrow \infty} \mathcal{P}^{(r)} \propto e^{-2\gamma^{(r)}t} \quad (4.20)$$

or

$$\lim_{t \rightarrow \infty} \frac{d}{dt} \mathcal{P}^{(r)} \sim -2\gamma^{(r)}t, \quad (4.21)$$

where γ is the decay rate of the slowest eigenmode of the system's hydrodynamics. We then identify the diffusion constant

$$D^{(r)} \equiv L^2\gamma^{(r)}/2\pi^2 \quad (4.22)$$

on the assumption that the slowest mode has characteristic wavevector π/L , even if it is not strictly the cosine one would find in the continuous diffusion equation with closed (homogeneous Neumann) boundary conditions and homogeneous diffusion coefficient. We then disorder-average¹

$$D = \left[\sum_{r=1}^{N_{\text{reals}}} (D^{(r)})^{-1} \right]^{-1} \quad (4.23)$$




Figure 4.6: Diffusion coefficients extracted from the distance from equilibrium (4.13). We perform a least-squares fit of each disorder realization’s distance from equilibrium $\log(\mathcal{P}^{(r)}(t))$ vs t for $924 \leq t \leq 1024$ for $L = 64, 128$ (or $150 \leq t \leq 250$ for $L = 16, 32$), and then disorder average per Eq. (4.23). The difference in times is because for $L = 16, 32$, by $t = 900$ the system is close enough to equilibrium that we do not believe we can reliably see deviations from that equilibrium. Error bars show 5th and 95th percentiles in a non-parametric bootstrap scheme.

Fig. 4.6 shows the disorder-averaged diffusion coefficients extracted from \mathcal{P}_4 at different disorder widths W . $L = 16$ shows slower diffusion than $L = 32, 64, 128$, but as noted above it does not properly reflect the long-system hydrodynamics. $L = 128$ shows somewhat faster diffusion at disorder widths $W \gtrsim 3.0$; this is because for $L \gtrsim 128, W \gtrsim 3.0$ the higher decay modes have not entirely decayed away by the end of our simulation time, so the diffusion coefficient we extract is polluted by that faster decay.

¹Note that in the language of Le Doussal and Machta [129] this is the so-called “annealed” diffusion coefficient, corresponding to an immediate spatial average, as opposed to the “quenched” diffusion coefficient. Some care must be taken in comparing our results here to those like [118] that probe different length scales at different times.




Figure 4.7: Another scaling collapse for the distance from equilibrium $\mathcal{P}(t)$ for 30 disorder realizations at $L = 64$. Error bars show 5th and 95th percentiles in a non-parametric bootstrap scheme.

These results, while suggestive, are not sufficient to demonstrate that our data shows diffusive behavior at *all* disorder widths $0 \leq W \leq 8$.² In Fig. 4.7 we plot the power $\mathcal{P}_4(t)$ against the rescaled time Dt for $L = 64$. All disorder widths $0 \leq W \leq 8$ show a good scaling collapse on the shape of figure 4.6. Since the diffusion constant is a free parameter in the scaling and we determine the diffusion coefficient from exactly and only the data shown in Fig. 4.7, there is a certain circularity in the analysis that is absent from the scaling collapse of Fig. 4.6; nonetheless, the consistency in the overall shape of the curve demonstrates that the system is in fact diffusive for $0 \leq W \leq 8$.

4.4 Discussion

We have applied time evolution with DMT to a disordered low-frequency Floquet system. We first evaluated its performance in this case both by checking convergence in bond dimension and by comparing with exact simulations. We found that it gives moderately good results, with error of order 5% in the largest exactly-simulable clean systems. But if one seeks greater precision than that, DMT may not be able to offer that precision. And, more worryingly, time evolution with DMT does not appear to be self-diagnosing with respect to errors of that order. We believe that the physics DMT fails to capture consists of finite size effects, perhaps peculiar to this model, but much work remains to be done to understand those effects and convincingly argue

²As the joke runs: we have thus far seen that at least one side of one sheep in Scotland is black.

that DMT will converge to better than this 5% precision in long systems.

Additionally, work remains to be done to reconcile the convergence tests of Sec. 4.2.1 with the comparison with exact results of Sec. 4.2.2. The comparison with exact results show that a bond dimension $\chi = 32$ gives a diffusion coefficient that differs from that of $\chi = 128$ by $\sim 20\%$ at $L = 28$, but the convergence tests at $L = 64$ show that local observables differ between $\chi = 32$ and $\chi = 128$ by $\lesssim 5\%$ in the clean case. This difference may be due to the different measures (diffusion coefficient vs. local observables), but more optimistically it could be due to improved performance of DMT at longer system sizes.

Additionally, it is possible that DMT fails to capture Griffiths regions—but we see apparent convergence because the systems we use to check convergence happen not to have rare, anomalously slow regions.

We found that—contra expectations based on [108, 147]—it displays a large regime where it is diffusive, not subdiffusive.

Part II

MBL-mobile: Many-body-localized engine

MBL AND THERMODYNAMIC ENGINES

In Ch. 4 we alluded to many body localization, but we did not treat a localized system. At most, many body localization could have caused regions of anomalously large resistance, leading to subdiffusion. But a tantalizing question is whether the unique properties of MBL could be put to use. So far, MBL was proposed to be used for robust quantum memories [158]. We believe, however, that the potential of MBL is much greater. MBL systems behave athermally, and athermality (lack of thermal equilibrium) facilitates thermodynamic tasks [8, 26, 43, 50, 55, 69, 75, 83, 91, 98, 229, 236]. When a cold bath is put in contact with a hot environment, for instance, work can be extracted from the heat flow. Could MBL's athermality have thermodynamic applications?

In this part we present such an application. We formulate, analyze, and numerically simulate an Otto engine cycle for a quantum many-body system that has an MBL phase. The engine contacts a hot bath and a narrow-band cold bath, as sketched in Fig. 5.1. This application unites the growing fields of quantum thermal machines [22, 28, 51, 65, 66, 99, 123, 130, 164, 199, 205, 232] and MBL [19, 92, 158–160, 191]. Our proposal could conceivably be explored in ultracold-atom [24, 37, 120, 148, 188], nitrogen-vacancy-center [125], trapped-ion [198], and possibly doped-semiconductor [124] experiments.

Our engine relies on two properties that distinguish MBL from thermal systems: its spectral correlations [191, 196] and its localization. The spectral-correlation properties enable us to build a mesoscale level-statistics engine. The localization enables us to link mesoscale engines together and create a large engine with extensive work output.

Take an interacting finite spin chain as an example. Consider the statistics of the gaps between consecutive energy eigenvalues far from the energy band's edges. A gap distribution $P(\delta)$ encodes the probability that any given gap has size δ . The MBL gap distribution enables small (and large) gaps to appear much more often than in ETH spectra [44]. This difference enables MBL to enhance our quantum many-body Otto cycle.

Let us introduce the MBL and ETH distributions in greater detail. Let $\langle \delta \rangle_E$




Figure 5.1: **Schematic of MBL engine:** We formulate an Otto engine cycle for a many-body quantum system that exhibits an MBL phase. We illustrate with a spin chain (green dots and black arrows). A random disorder potential (jagged red line) localizes the particles. Particles interact and hop between sites (horizontal red arrows). Consider strengthening the interactions and the hopping frequency. The system transitions from strong localization to a thermal phase or to weak localization. The engine thermalizes with a hot bath (flames) and with a cold bath (ice cube). The cold bath has a small bandwidth W_b , to take advantage of small energy gaps' greater prevalence in the highly localized regime.

denote the average gap at the energy E . MBL gaps approximately obey Poisson statistics [44, 159]:

$$P_{\text{MBL}}^{(E)}(\delta) \approx \frac{1}{\langle \delta \rangle_E} e^{-\delta / \langle \delta \rangle_E}. \quad (5.1)$$

Any given gap has a decent chance of being small: As $\delta \rightarrow 0$, $P_{\text{MBL}}^{(E)}(\delta) \rightarrow \frac{1}{\langle \delta \rangle_E} > 0$. Neighboring energies have finite probabilities of lying close together: MBL systems' energies do not repel each other, unlike thermal systems' energies. Thermalizing systems governed by real Hamiltonians obey the level statistics of random matrices drawn from the Gaussian orthogonal ensemble (GOE) [159]:

$$P_{\text{GOE}}^{(E)}(\delta) \approx \frac{\pi}{2} \frac{\delta}{\langle \delta \rangle_E^2} e^{-\frac{\pi}{4} \delta^2 / \langle \delta \rangle_E^2}. \quad (5.2)$$

Unlike in MBL spectra, small gaps rarely appear: As $\delta \rightarrow 0$, $P_{\text{GOE}}^{(E)}(\delta) \rightarrow 0$.

MBL's athermal gap statistics should be construed as a thermodynamic resource, we find, as athermal quantum states are [8, 26, 43, 50, 55, 69, 75, 83, 91, 98, 229, 236]. In particular, MBL's athermal gap statistics improve our engine's reliability: The amount W_{tot} of work extracted by our engine fluctuates relatively little from successful trial to successful trial. Athermal statistics

also lower the probability of worst-case trials, in which the engine outputs net negative work, $W_{\text{tot}} < 0$. Furthermore, MBL’s localization enables the engine to scale robustly: Mesoscale “subengines” can run in parallel without disturbing each other much, due to the localization inherent in MBL. Even in the thermodynamic limit, an MBL system behaves like an ensemble of finite, mesoscale quantum systems, due to its *local level correlations* [96, 196, 201]. Any local operator can probe only a discrete set of sharp energy levels, which emerge from its direct environment.

This part is organized as follows. Ch. 6 describes the MBL-mobile in the adiabatic limit—we take the tuning to be infinitely slow. In Sec. 6.1, we introduce the basic idea with a qubit (two-level quantum system). In Sec. 6.2, we scale the engine up to a mesoscopic chain tuned between MBL and ETH. In Sec. 6.5, we calculate its work output and efficiency. We check these calculations numerically in Sec. 6.7 with numerical simulations of disordered spin chains. In Sec. 6.6, we argue that the mesoscopic segments can be combined into a macroscopic MBL system while operating in parallel.

Ch. 7 describes the MBL-mobile at finite speed. We discuss limitations on the speed at which the engine can be run (and consequently the engine’s power). This leads us to a more careful consideration of diabatic corrections to the work output, communication amongst subengines, and the cold bath’s nature. We again test our analytic calculations in Sec. 7.2,

In Ch. 8 we provide order-of-magnitude estimates for a localized semiconductor engine’s power and power density, and we compare our engine to some competitor engines.

Background information, intuitive examples, and extensive calculations appear in [238].

THE MBL-MOBILE IN ITS ADIABATIC LIMIT

We aim to formulate an MBL engine cycle for the thermodynamic limit. Our road to that goal runs through a finite-size, or mesoscale, MBL engine. In Sec. 6.1, we introduce the intuition behind the mesoscale engine via a qubit toy model. Then, we describe (Sec. 6.2) and quantitatively analyze (Sec. 6.5) the mesoscale MBL engine. Table 6.1 offers a spotter’s guide to notation.

6.1 Qubit toy model for the mesoscale engine

At the MBL Otto engine’s core lies a qubit Otto engine whose energy eigenbasis transforms during the cycle [31, 112, 121, 122]. Consider a two-level system evolving under the time-varying Hamiltonian

$$H_{\text{qubit}}(t) := (1 - \alpha_t)h\sigma^x + \alpha_t h'\sigma^z. \quad (6.1)$$

σ^x and σ^z denote the Pauli x - and z -operators. α_t denotes a parameter tuned between 0 and 1.

Figure 6.1 illustrates the cycle. The engine begins in thermal equilibrium at a high temperature T_H . During stroke 1, the engine is thermally isolated, and α_t is tuned from 0 to 1. During stroke 2, the engine thermalizes to a temperature $T_C \ll T_H$. During stroke 3, the engine is thermally isolated, and α_t returns from 1 to 0. During stroke 4, the engine resets by thermalizing with the hot bath.

Let us make two simplifying assumptions (see [238, App. C] for a generalization): First, let $T_H = \infty$ and $T_C = 0$. Second, assume that the engine is tuned slowly enough to satisfy the quantum adiabatic theorem. We also choose¹

$$h = \frac{\delta_{\text{GOE}}}{2}, \quad h' = \frac{\delta_{\text{MBL}}}{2},$$

and $\delta_{\text{GOE}} \gg \delta_{\text{MBL}}$.

¹The gaps’ labels are suggestive: A qubit, having only one gap, obeys neither GOE nor MBL gap statistics. But, when large, the qubit gap apes a typical GOE gap; and, when small, the qubit gap apes a useful MBL gap. This mimicry illustrates how the mesoscopic engine benefits from the greater prevalence of small gaps in MBL spectra than in GOE spectra.




Figure 6.1: **Qubit toy model for the MBL Otto cycle:** A qubit models two “working levels” in the MBL Otto engine’s many-body spectrum. The energy eigenstates $|E_t^{(1)}\rangle$ and $|E_t^{(2)}\rangle$ span the “working subspace.” The gap $E_t^{(2)} - E_t^{(1)}$ begins at size δ_{GOE} during a successful trial. The gap shrinks to δ_{MBL} , then returns to δ_{GOE} . In addition to changing the gap, each Hamiltonian tuning changes the eigenstates’ functional forms. The displacement δ_{displ} is included for generality. The blue text marks the times $t = 0, \tau, \dots, \tau'''$ at which the strokes begin and end during a work-outputting trial. The spectator level $|E_t^{(3)}\rangle$ fails to impact the engine’s efficiency. The cold bath has too narrow a bandwidth W_b to couple $|E_t^{(3)}\rangle$ to any other level. If the engine begins any trial on the top green line, the engine remains on that line throughout the trial. Zero net work is outputted.

Let us analyze the cycle’s energetics. The system begins with $\langle H_{\text{qubit}}(t) \rangle = 0$. Stroke 1 preserves the $T = \infty$ state $I/2$. Stroke 2 drops the energy to $-\frac{\delta_{\text{MBL}}}{2}$. The energy drops to $-\frac{\delta_{\text{GOE}}}{2}$ during stroke 3. During stroke 4, the engine resets to zero average energy, absorbing heat $\langle Q_4 \rangle = \frac{\delta_{\text{GOE}}}{2}$, on average.

The energy exchanged during the tunings (strokes 1 and 3) constitutes work [Eq. (1.137)], while the energy exchanged during the thermalizations (strokes 2 and 4) is heat [Eq. (1.138)]. The engine outputs the *per-cycle power*, or average work performed per cycle, $\langle W_{\text{tot}} \rangle = \frac{1}{2}(\delta_{\text{GOE}} - \delta_{\text{MBL}})$.

The efficiency is $\eta_{\text{qubit}} = \frac{\langle W_{\text{tot}} \rangle}{\langle Q_4 \rangle} = 1 - \frac{\delta_{\text{MBL}}}{\delta_{\text{GOE}}}$. This result is equivalent to the




Figure 6.2: **Otto engine cycle for a mesoscale MBL system:** Two energies in the many-body spectrum capture the cycle’s basic physics. The engine can be regarded as beginning each trial in an energy eigenstate drawn from a Gibbs distribution. The red dot represents the engine’s starting state in some trial of interest. During stroke 1, $H_{\text{meso}}(t)$ is tuned from “thermal” to MBL. During stroke 2, the engine thermalizes with a cold bath. $H_{\text{meso}}(t)$ returns from MBL to thermal during stroke 3. Stroke 4 resets the engine, which thermalizes with a hot bath. The tunings (strokes 1 and 3) map onto the thermodynamic Otto cycle’s isentropes. The thermalizations (strokes 2 and 4) map onto isochores. The engine outputs work W_1 and W_3 during the tunings and absorbs heat Q_2 and Q_4 during thermalizations. MBL gap statistics’ lack of level repulsion enhances the cycle: The engine “slides down” the lines that represent tunings, losing energy outputted as work.

efficiency η_{Otto} of a thermodynamic Otto engine [Eq. (1.135)]. The gap ratio $\frac{\delta_{\text{MBL}}}{\delta_{\text{GOE}}}$ plays the role of $\frac{1}{r^{\gamma-1}}$. η_{qubit} equals also η_{QHO} [Eq. (1.136)] if the frequency ratio ω/Ω is chosen to equal $\delta_{\text{MBL}}/\delta_{\text{GOE}}$. As argued in Sec. 6.6 (and 7.3), however, the qubit engine can scale to a large composite engine of densely packed qubit subengines operating in parallel. The dense packing is possible if the qubits are encoded in the MBL system’s localized degrees of freedom (1-bits, roughly speaking [93]).

6.2 Set-up for the mesoscale MBL engine

The next step is an interacting finite-size system tuned between MBL and ETH phases. Envision a mesoscale engine as a one-dimensional (1D) system of $N \approx$

Symbol	Significance
N	Number of sites per mesoscale engine or per mesoscale subengine (in the macroscopic engine). Chosen, in the latter case, to equal $\xi_{>}$.
\mathcal{N}	Dimensionality of one mesoscale (sub)engine's Hilbert space.
\mathcal{E}	Unit of energy, average energy density per site.
α_t	Hamiltonian parameter tuned from 0 (in the mesoscale engine's ETH regime, or the macroscopic engine's shallowly localized regime) to 1 (in the engine's deeply MBL regime).
$\langle\delta\rangle$	Average gap in the energy spectrum of a length- N MBL system.
W_b	Bandwidth of the cold bath. Small: $W_b \ll \langle\delta\rangle$.
$\beta_H = 1/T_H$	Inverse temperature of the hot bath.
$\beta_C = 1/T_C$	Inverse temperature of the cold bath.
δ_-	Level-repulsion scale of a length- N MBL system. Minimal size reasonably attributable to any energy gap. Smallest gap size at which a Poissonian (5.1) approximates the MBL gap distribution well.
v	Speed at which the Hamiltonian is tuned: $v := \mathcal{E} \frac{d\alpha_t}{dt}$. Has dimensions of 1/time ² , in accordance with part of [46].
$\xi_{>}$	Localization length of macroscopic MBL engine when shallowly localized. Length of mesoscale subengine.
$\xi_{<}$	Localization length of macroscopic MBL engine when deeply localized. Satisfies $\xi_{<} < \xi_{>}$.
X_{macro}	Characteristic X of the macroscopic MBL engine (e.g., $X = N, \langle\delta\rangle$).
g	Strength of coupling between engine and cold bath.
τ_{cycle}	Time required to implement one cycle.
$\langle\delta\rangle^{(L)}$	Average energy gap of a length- L MBL system.

Table 6.1: **Parameters of the mesoscopic and macroscopic MBL engines:** Boltzmann's and Planck's constants are set to one: $k_B = \hbar = 1$.

10 sites. This engine will ultimately model one region in a thermodynamically large MBL engine. We will analyze the mesoscopic engine's per-trial power $\langle W_{\text{tot}} \rangle$, the efficiency η_{MBL} , and work costs $\langle W_{\text{diab}} \rangle$ of undesirable diabatic transitions.

The mesoscopic engine evolves under the Hamiltonian

$$H_{\text{meso}}(t) := \frac{\mathcal{E}}{Q(\alpha_t)} [(1 - \alpha_t)H_{\text{GOE}} + \alpha_t H_{\text{MBL}}]. \quad (6.2)$$

The unit of energy, or average energy density per site, is denoted by \mathcal{E} . The tuning parameter $\alpha_t \in [0, 1]$. When $\alpha_t = 0$, the system evolves under a random Hamiltonian H_{GOE} whose gaps δ are distributed according to $P_{\text{GOE}}^{(E)}(\delta)$ [Eq. (5.2)]. When $\alpha_t = 1$, $H_{\text{meso}}(t) = H_{\text{MBL}}$, a Hamiltonian whose gaps are distributed according to $P_{\text{MBL}}^{(E)}(\delta)$ [Eq. (5.1)]. For a concrete example, take a random-field Heisenberg model in which the disorder strength is tuned. H_{GOE} and H_{MBL} have the same bond term, but the disorder strength varies in time. We simulate (a rescaled version of) this model in Sec. 6.7.

The mesoscale engine’s cycle is analogous to the qubit cycle, including initialization at $\alpha_t = 0$, tuning of α_t to one, thermalization with a temperature- T_C bath, tuning of α_t to zero, and thermalization [47, 63, 94, 132] with a temperature- T_H bath. To highlight the role of level statistics in the cycle, we hold the average energy gap $\langle \delta \rangle$ constant.² We do so using the renormalization factor $Q(\alpha_t)$.³ Section 6.7 details how we define $Q(\alpha_t)$ in numerical simulations.

The key distinction between GOE level statistics (5.2) and Poisson (MBL) statistics (5.1) is that small gaps (and large gaps) appear more often in Poisson spectra. A toy model illuminates these level statistics’ physical origin: An MBL system can be modeled as a set of noninteracting quasilocal qubits [93]. Let g_j denote the j th qubit’s gap. Two qubits, j and j' , may have nearly equal gaps: $g_j \approx g_{j'}$. The difference $|g_j - g_{j'}|$ equals a gap in the many-body energy spectrum. Tuning the Hamiltonian from MBL to ETH couples

² $\langle \delta \rangle$ is defined as follows. The density of states at the energy E has the form $\mu(E) \approx \frac{\mathcal{N}}{\sqrt{2\pi N} \mathcal{E}} e^{-E^2/2N\mathcal{E}^2}$ (see Table 6.1 for the symbols’ meanings). Inverting $\mu(E)$ yields the *local average gap*: $\langle \delta \rangle_E := \frac{1}{\mu(E)}$. Inverting the average of $\mu(E)$ yields the *average gap*,

$$\langle \delta \rangle := \frac{1}{\langle \mu(E) \rangle_{\text{energies}}} = \frac{\mathcal{N}}{\int_{-\infty}^{\infty} dE \mu^2(E)} = \frac{2\sqrt{\pi N}}{\mathcal{N}} \mathcal{E}. \quad (6.3)$$

³Imagine removing $Q(\alpha_t)$ from Eq. (6.2). One could increase α_t —could tune the Hamiltonian from ETH to MBL [191]—by strengthening a disorder potential. This strengthening would expand the energy band; tuning oppositely would compress the band. By expanding and compressing, in accordion fashion, and thermalizing, one could extract work. This engine would benefit little from properties of MBL, whose thermodynamic benefits we wish to highlight. So we “zero out” the accordion motion by fixing $\langle \delta \rangle$ through $Q(\alpha_t)$. For a brief discussion of the accordion-like engine, see Sec. 8.1.1.

the qubits together, producing matrix elements between the nearly degenerate states. These matrix elements force energies apart.

To take advantage of the phases' distinct level statistics, we use a cold bath that has a small bandwidth W_b . According to Sec. 6.1, net positive work is extracted from the qubit engine because $\delta_{\text{MBL}} < \delta_{\text{GOE}}$. The mesoscale analog of δ_{GOE} is $\sim \langle \delta \rangle$, the typical gap ascended during hot thermalization. The engine must not emit energy on this scale during cold thermalization. Limiting W_b ensures that cold thermalization relaxes the engine only across gaps $\delta \leq W_b \ll \langle \delta \rangle$. Such anomalously small gaps appear more often in MBL energy spectra than in ETH spectra [71, 107, 161].

This level-statistics argument holds only within superselection sectors. Suppose, for example, that $H_{\text{meso}}(t)$ conserves particle number. The level-statistics arguments apply only if the particle number remains constant throughout the cycle [238, App. F]. Our numerical simulations (Sec. 6.7) take place at half-filling, in a subspace of dimensionality \mathcal{N} of the order of magnitude of the whole space's dimensionality: $\mathcal{N} \sim \frac{2^N}{\sqrt{N}}$.

We are now ready to begin analyzing the mesoscopic-engine Otto cycle. The engine begins in the thermal state $\rho(0) = e^{-\beta_{\text{H}} H_{\text{GOE}}} / Z$, wherein $Z := \text{Tr} (e^{-\beta_{\text{H}} H_{\text{GOE}}})$. The engine can be regarded as starting each trial in some energy eigenstate j drawn according to the Gibbs distribution (Fig. 6.2). During stroke 1, $H_{\text{meso}}(t)$ is tuned from H_{GOE} to H_{MBL} . We approximate the tuning as quantum-adiabatic (diabatic corrections are left to Sec. 7.1). Stroke 2, cold thermalization, depends on the gap δ'_j between the j th and $(j-1)$ th MBL levels. δ'_j typically exceeds W_b . If it does, cold thermalization preserves the engine's energy, and the cycle outputs $W_{\text{tot}} = 0$. With probability $\sim \frac{W_b}{\langle \delta \rangle}$, the gap is small enough to thermalize: $\delta'_j < W_b$. In this case, cold thermalization drops the engine to level $j-1$. Stroke 3 brings the engine to level $j-1$ of H_{GOE} . The gap δ_j between the $(j-1)$ th and j th H_{GOE} levels is $\langle \delta \rangle \gg W_b$, with the high probability $\sim 1 - (W_b / \langle \delta \rangle)^2$. Therefore the engine likely outputs $W_{\text{tot}} > 0$. Hot thermalization (stroke 4) returns the engine to $\rho(0)$.

6.3 Notation and definitions:

We focus on one mesoscopic engine S of N sites. The engine corresponds to a Hilbert space of dimensionality $\mathcal{N} \sim \frac{2^N}{\sqrt{N}}$. The Hamiltonian, $H(t) \equiv H_{\text{meso}}(t)$, is tuned between H_{GOE} , which obeys the ETH, and H_{MBL} , which governs an

MBL system. ETH and MBL Hamiltonians typically have Gaussian DoSs⁴

$$\mu(E) = \frac{\mathcal{N}}{\sqrt{2\pi N} \mathcal{E}} e^{-E^2/(2N\mathcal{E}^2)}, \quad (6.4)$$

normalized to $\int_{-\infty}^{\infty} dE \mu(E) = \mathcal{N}$. The unit of energy, or energy density per site, is \mathcal{E} . We often extend energy integrals' limits to $\pm\infty$, as the Gaussian peaks sharply about $E = 0$.

The local average gap is $\langle\delta\rangle_E = \frac{1}{\mu(E)}$, and the average gap is $\langle\delta\rangle := \frac{\mathcal{N}}{\int_{-\infty}^{\infty} dE \mu^2(E)} = \frac{2\sqrt{\pi N} \mathcal{E}}{\mathcal{N}}$ [Eq. 6.3]. The average H_{GOE} gap, $\langle\delta\rangle$, equals the average H_{MBL} gap, by construction. $\langle\delta\rangle$ sets the scale for work and heat quantities. Hence we cast Q 's and W 's as (number)(function of small parameters) $\langle\delta\rangle$.

The system begins the cycle in the state $\rho(0) = e^{-\beta_{\text{H}} H_{\text{GOE}}}/Z$, wherein $Z := \text{Tr}(e^{-\beta_{\text{H}} H_{\text{GOE}}})$ denotes the partition function. W_{b} denotes the cold bath's bandwidth. We set $\hbar = k_{\text{B}} = 1$.

$H(t)$ is tuned at a speed $v := \mathcal{E} \left| \frac{d\alpha_t}{dt} \right|$, wherein α_t denotes the dimensionless tuning parameter. v has dimensions of energy², as in [195]. Though our v is not defined identically to the v in [195], we expect ours to behave similarly.

6.4 Quantitative analysis of the adiabatic mesoscale engine: some easy limits

How well does the mesoscale Otto engine perform? We calculate average work $\langle W_{\text{tot}} \rangle$ outputted per cycle and the efficiency η_{MBL} .

We focus on the parameter regime in which the cold bath is very cold, the cold-bath bandwidth W_{b} is very small, and the hot bath is very hot: $T_{\text{C}} \ll W_{\text{b}} \ll \langle\delta\rangle$, and $\sqrt{N} \beta_{\text{H}} \mathcal{E} \ll 1$. The mesoscale engine resembles a qubit engine whose state and gaps are averaged over. The gaps, δ_j and δ'_j , obey the distributions $P_{\text{GOE}}^{(E)}(\delta_j)$ and $P_{\text{MBL}}^{(E)}(\delta'_j)$ [Eqs. (5.2) and (5.1)]. Correlations between the H_{GOE} and H_{MBL} spectra can be neglected.

In this section we make three simplifying assumptions: (i) The engine is as-

⁴This is only approximately true, even in the thermodynamic limit. In the random-field Heisenberg Hamiltonian (6.33) that we use for our numerical simulations, for example, the coupling term skews the the density of states: the mean of the DoS is $\text{tr} H = 0$, but the median is nonzero. Our numerics nonetheless match analytics performed assuming a Gaussian DoS (cf Sec. 6.7). This is because the Gaussian is sharply peaked around the mode: we can think of the skewness as an effective constant offset to the energy.

Moreover, the DoS is essentially a parameter in our calculations: the details of some of our results are specific to this form for the DoS, but the physics does not depend on it.

sumed to be tuned quantum-adiabatically. (ii) The hot bath is at $T_H = \infty$. We neglect finite-temperature corrections, which scale as $N(\beta_H \mathcal{E})^2 \frac{W_b}{\langle \delta \rangle^2}$. (iii) The gap distributions vary negligibly with energy: $P_{\text{GOE}}^{(E)}(\delta_j) \approx P_{\text{GOE}}(\delta_j)$, and $P_{\text{MBL}}^{(E)}(\delta'_j) \approx P_{\text{MBL}}(\delta'_j)$, wherein $\langle \delta \rangle_E \approx \langle \delta \rangle$.

We relax assumptions (ii) and (iii) in Sec. 6.4; relaxing assumption (i) must wait till Sec. 7.1.

6.4.0.1 Average work $\langle W_{\text{tot}} \rangle$ per cycle:

The key is whether the cold bath relaxes the engine downwards across the MBL-side gap $\delta' \equiv \delta'_j$, distributed as $P_{\text{MBL}}(\delta')$, during a given trial. If $\delta' < W_b$, the engine has a probability $1/(1 + e^{-\beta_C \delta'})$ of thermalizing. So the overall probability of relaxation by the cold bath is

$$p_{\text{cold}} \approx \int_0^{W_b} d\delta' \frac{1}{\langle \delta \rangle} \frac{e^{-\delta'/\langle \delta \rangle}}{1 + e^{-\beta_C \delta'}} \approx \frac{1}{\langle \delta \rangle} (W_b - T_C \ln 2), \quad (6.5)$$

wherein we neglected $W_b/\langle \delta \rangle$ by setting $e^{-\delta'/\langle \delta \rangle} \approx 1$.

Alternatively, the cold bath could excite the engine to a level a distance δ' above the initial level. Such an upward hop occurs with a probability

$$\bar{p}_{\text{cold}} \approx \int_0^{W_b} d\delta' \frac{e^{-\delta'/\langle \delta \rangle}}{\langle \delta \rangle} \frac{e^{-\beta_C \delta'}}{1 + e^{-\beta_C \delta'}} \approx \frac{T_C \ln 2}{\langle \delta \rangle}. \quad (6.6)$$

If the engine relaxed downward during stroke 2, then upon thermalizing with the hot bath during stroke 4, the engine gains heat $\langle Q \rangle_4 \approx \langle \delta \rangle$, on average. If the engine thermalized upward during stroke 2, then the engine loses $\langle \delta \rangle$ during stroke 4, on average. Therefore, the cycle outputs average work

$$\langle W_{\text{tot}} \rangle \approx (p_{\text{cold}} - \bar{p}_{\text{cold}}) \langle \delta \rangle + \langle Q_2 \rangle \approx W_b - \frac{2 \ln 2}{\beta_C}. \quad (6.7)$$

$\langle Q_2 \rangle$ denotes the average heat absorbed by the engine during cold thermalization:

$$\langle Q_2 \rangle \approx - \int_0^{W_b} d\delta' \frac{\delta'}{\langle \delta \rangle} \frac{e^{-\delta'/\langle \delta \rangle}}{1 + e^{-\beta_C \delta'}} \approx - \frac{(W_b)^2}{2 \langle \delta \rangle}, \quad (6.8)$$

which is $\ll \langle Q_4 \rangle$. This per-cycle power scales with the system size N as⁵ $W_b \ll \langle \delta \rangle \sim \frac{\text{effective bandwidth}}{\# \text{ energy eigenstates}} \sim \frac{\mathcal{E}\sqrt{N}}{N}$.

6.4.0.2 Efficiency η_{MBL} :

The efficiency is

$$\eta_{\text{MBL}} = \frac{\langle W_{\text{tot}} \rangle}{\langle Q_4 \rangle} = \frac{\langle Q_4 \rangle + \langle Q_2 \rangle}{\langle Q_4 \rangle} \approx 1 - \frac{W_b}{2\langle \delta \rangle}. \quad (6.9)$$

The imperfection is small, $\frac{W_b}{2\langle \delta \rangle} \ll 1$, because the cold bath has a small bandwidth. This result mirrors the qubit-engine efficiency η_{qubit} .⁶ But our engine is a many-body system of N interacting sites. MBL will allow us to employ segments of the system as independent qubit-like subengines, despite interactions. In the absence of MBL, each subengine's effective $\langle \delta \rangle = 0$. With $\langle \delta \rangle$ vanishes the ability to extract $\langle W_{\text{tot}} \rangle > 0$. Whereas the efficiency is nearly perfect, an effective engine requires also optimized power. The MBL engine's power will be limited by dynamical considerations, discussed in Ch. 7.

6.5 Quantitative analysis of the adiabatic mesoscale engine: a more detailed calculation

6.5.1 Partial-swap model of thermalization

Classical thermalization can be modeled with a *probabilistic swap*, or *partial swap*, or *p-SWAP* [186, 245]. Let a column vector \vec{v} represent the state. The thermalization is broken into time steps. At each step, a doubly stochastic matrix M_p operates on \vec{v} . The matrix's fixed point is a Gibbs state \vec{g} .

M_p models a probabilistic swapping out of \vec{v} for \vec{g} : At each time step, the system's state has a probability $1 - p$ of being preserved and a probability

⁵The *effective bandwidth* is defined as follows. The many-body system has a Gaussian density of states: $\mu(E) \approx \frac{\mathcal{N}}{\sqrt{2\pi N} \mathcal{E}} e^{-E^2/2N\mathcal{E}^2}$. The states within a standard deviation $\mathcal{E}\sqrt{N}$ of the mean have gaps obeying the Poisson or GOE level-spacing distributions, Eqs. (5.1) and (5.2). These states form the effective band, whose width scales as $\mathcal{E}\sqrt{N}$.

⁶ η_{MBL} is comparable also to η_{QHO} [Eq. (1.136)]. Imagine operating an ensemble of independent QHO engines. Let the j th QHO frequency be tuned between Ω_j and ω_j , distributed according to $P_{\text{GOE}}(\Omega_j)$ and $P_{\text{MBL}}(\omega_j)$. The average MBL-like gap ω_j , conditioned on $\omega_j \in [0, W_b]$, is $\langle \omega_j \rangle \sim \frac{1}{W_b/\langle \delta \rangle} \int_0^{W_b} d\omega_j \omega_j P_{\text{MBL}}(\omega_j) \approx \frac{1}{W_b} \int_0^{W_b} d\omega_j \omega_j = \frac{W_b}{2}$. Averaging the efficiency over the QHO ensemble yields $\langle \eta_{\text{QHO}} \rangle := 1 - \frac{\langle \omega \rangle}{\langle \Omega \rangle} \approx 1 - \frac{W_b}{2\langle \delta \rangle} \approx \eta_{\text{MBL}}$. The mesoscale MBL engine operates at the ideal average efficiency of an ensemble of QHO engines. But MBL enables qubit-like engines to pack together densely in a large composite engine.

$p \in [0, 1]$ of being replaced by \vec{g} . This algorithm gives M_p the form $M_p = (1 - p)I + p\vec{g}(1, 1)$.

We illustrate with thermalization across two levels. Let 0 and Δ label the levels, such that $\vec{g} = \left(\frac{e^{-\beta\Delta}}{1+e^{-\beta\Delta}}, \frac{1}{1+e^{-\beta\Delta}} \right)$:

$$M_p = \begin{bmatrix} 1 - p \frac{1}{1+e^{-\beta\Delta}} & p \frac{e^{-\beta\Delta}}{1+e^{-\beta\Delta}} \\ p \frac{1}{1+e^{-\beta\Delta}} & 1 - p \frac{e^{-\beta\Delta}}{1+e^{-\beta\Delta}} \end{bmatrix}. \quad (6.10)$$

The off-diagonal elements, or transition probabilities, obey detailed balance [40, 239]: $\frac{P(0 \rightarrow \Delta)}{P(\Delta \rightarrow 0)} = e^{-\beta\Delta}$.

Repeated application of M_p maps every state to \vec{g} [239]: $\lim_{n \rightarrow \infty} (M_p)^n \vec{v} = \vec{g}$. The parameter p reflects the system-bath-coupling strength. We choose $p = 1$: The system thermalizes completely at each time step. (If $p \neq 1$, a more sophisticated model may be needed for thermalization across > 2 levels.)

6.5.2 Average heat $\langle Q_2 \rangle$ absorbed during stroke 2

Let j denote the H_{GOE} level in which the engine begins the trial of interest. We denote by $Q_2^{(j)}$ the average heat absorbed during stroke 2, from the cold bath. ($Q_2^{(j)}$ will be negative and, provided that j is around the energy band's center, independent of j .)

The heat absorbed can be calculated easily from the following observation. Stroke 1 (adiabatic tuning) preserves the occupied level's index. The level closest to j lies a distance δ away when stroke 3 begins. δ can have either sign, can lie above or below j . Heat is exchanged only if $|\delta| < W_b$. Let us initially neglect the possibility that two nearby consecutive gaps are very small: $|E_{j\pm 2} - E_j| \leq W_b$. We can write the average (over trials begun in level j) heat absorbed as

$$Q_2^{(j)} = \int_{-W_b}^{W_b} d\delta \delta \frac{e^{-\beta_C \delta}}{1 + e^{-\beta_C \delta}} P_{\text{MBL}}^{(E)}(\delta) + O(W_b^3 / \langle \delta \rangle^2). \quad (6.11)$$

This equation assumes a Sommerfeld-expansion form, so

$$Q_2^{(j)} = -\frac{W_b^2}{2} \mu(E) + \frac{\pi^2}{6} \mu(E) (T_C)^2 + O([W_b]^3 / \langle \delta \rangle^2) + O(\mu(E)^2 [T_C]^3). \quad (6.12)$$

The first correction accounts for our not considering two levels within W_b of level j .

Next, we need to average this result over all initial states j , assuming the initial density operator, $\rho(0) = e^{-\beta_H H_{\text{GOE}}}/Z$:

$$\langle Q_2 \rangle := \left\langle \left\langle \langle Q_2(E) \rangle_{\text{therm.}}^{\text{cold}} \right\rangle_{\text{gaps}} \right\rangle_{\rho(0)} \quad (6.13)$$

$$= \left(-\frac{(W_b)^2}{2} + \frac{\pi^2}{6} \frac{1}{(\beta_C)^2} \right) \int_{-\infty}^{\infty} dE \mu^2(E) \frac{e^{-\beta_H E}}{Z} \quad (6.14)$$

$$+ \langle \delta \rangle \left\{ O \left(\left[\frac{W_b}{\langle \delta \rangle} \right]^3 \right) + O \left(\frac{W_b}{\langle \delta \rangle} e^{-\beta_C W_b} \right) + O \left(\left[\frac{\mu(E)}{\beta_C} \right]^3 \right) \right\}. \quad (6.15)$$

We substitute in for the DoS from Eq. (6.4):

$$\langle Q_2 \rangle = \frac{\mathcal{N}^2}{2\pi N \mathcal{E}^2} \frac{1}{Z} \left(-\frac{(W_b)^2}{2} + \frac{\pi^2}{6} \frac{1}{(\beta_C)^2} \right) \int_{-\infty}^{\infty} dE e^{-E^2/N\mathcal{E}^2} e^{-\beta_H E} + O(\cdot), \quad (6.16)$$

wherein the correction terms are abbreviated. The integral evaluates to $\sqrt{\pi N} \mathcal{E} e^{N(\beta_H \mathcal{E})^2/4}$.

The partition function is

$$Z = \int_{-\infty}^{\infty} dE \mu(E) e^{-\beta_H E} = \mathcal{N} e^{N(\beta_H \mathcal{E})^2/2}. \quad (6.17)$$

Substituting into Eq. (6.16) yields

$$\begin{aligned} \langle Q_2 \rangle = & \left(-\frac{(W_b)^2}{2 \langle \delta \rangle} + \frac{\pi^2}{6} \frac{1}{(\beta_C)^2 \langle \delta \rangle} \right) e^{-N(\beta_H \mathcal{E})^2/4} \\ & + \langle \delta \rangle \left\{ O \left(\left[\frac{W_b}{\langle \delta \rangle} \right]^3 \right) + O \left([\mu(E) W_b] \frac{\mu(E)}{\beta_C} e^{-\beta_C W_b} \right) \right. \\ & \left. + O \left(\left[\frac{\mu(E)}{\beta_C} \right]^3 \right) + O \left([\sqrt{N} \beta_H \mathcal{E}]^4 \right) \right\}. \end{aligned} \quad (6.18)$$

The prefactor was replaced with $\frac{1}{\langle \delta \rangle}$ via Eq. (6.3) in the main text.

Equation (6.18) is compared with numerical simulations in Fig. 6.3. In the appropriate regime (wherein $W_b \ll \langle \delta \rangle$ and $T_C \ll W_b$), the analytics agree well with the numerics, to within finite-size effects.

In terms of small dimensionless parameters,

$$\langle Q_2 \rangle = \langle \delta \rangle \left[-\frac{1}{2} \left(\frac{W_b}{\langle \delta \rangle} \right)^2 + \frac{\pi^2}{6} \frac{1}{(\beta_C \langle \delta \rangle)^2} \right] \left[1 - \frac{N}{4} (\beta_H \mathcal{E})^2 \right] + O(\cdot). \quad (6.19)$$

The leading-order term is second-order. So is the β_C correction; but $\frac{1}{(\beta_C \langle \delta \rangle)^2} \ll \left(\frac{W_b}{\langle \delta \rangle} \right)^2$, by assumption. The β_H correction is fourth-order—too small to include. To lowest order,

$$\boxed{\langle Q_2 \rangle \approx -\frac{(W_b)^2}{2 \langle \delta \rangle}}. \quad (6.20)$$

6.5.3 Average heat $\langle Q_4 \rangle$ absorbed during stroke 4

The $\langle Q_4 \rangle$ calculation proceeds similarly to the $\langle Q_2 \rangle$ calculation. When calculating $\langle Q_2 \rangle$, however, we neglected contributions from the engine's cold-thermalizing down two small gaps. Two successive gaps have a probability $\sim \left(\frac{W_b}{\langle \delta \rangle} \right)^2$ of being $< W_b$ each. Thermalizing across each gap produces heat $\leq W_b$. Each such pair therefore contributes negligibly to $\langle Q_2 \rangle$, as $\langle \delta \rangle O\left(\left[\frac{W_b}{\langle \delta \rangle}\right]^3\right)$.

We cannot neglect these pairs when calculating $\langle Q_4 \rangle$. Each typical small gap widens, during stroke 3, to size $\sim \langle \delta \rangle$. These larger gaps are thermalized across during stroke 4, contributing at the nonnegligible second order, as $\sim \langle \delta \rangle O\left(\left[\frac{W_b}{\langle \delta \rangle}\right]^2\right)$ to $\langle Q_4 \rangle$. Chains of ≥ 3 small MBL gaps contribute negligibly.

The calculation is tedious, appears in [238, App. G 5], and yields

$$\boxed{\langle Q_4 \rangle \approx W_b - \frac{2 \ln 2}{\beta_C} + \frac{(W_b)^2}{2 \langle \delta \rangle} + 4 \ln 2 \frac{W_b}{\beta_C \langle \delta \rangle}}. \quad (6.21)$$

The leading-order terms are explained heuristically below Eq. (6.7) in the main text.

The leading-order β_C correction, $-\frac{2 \ln 2}{\beta_C}$, shows that a warm cold bath lowers the heat required to reset the engine. Suppose that the cold bath is maximally cold: $T_C = 0$. Consider any trial that S begins just above a working gap (an ETH gap $\delta > W_b$ that narrows to an MBL gap $\delta' < W_b$). Cold thermalization drops S deterministically to the lower level. During stroke 4, S must absorb

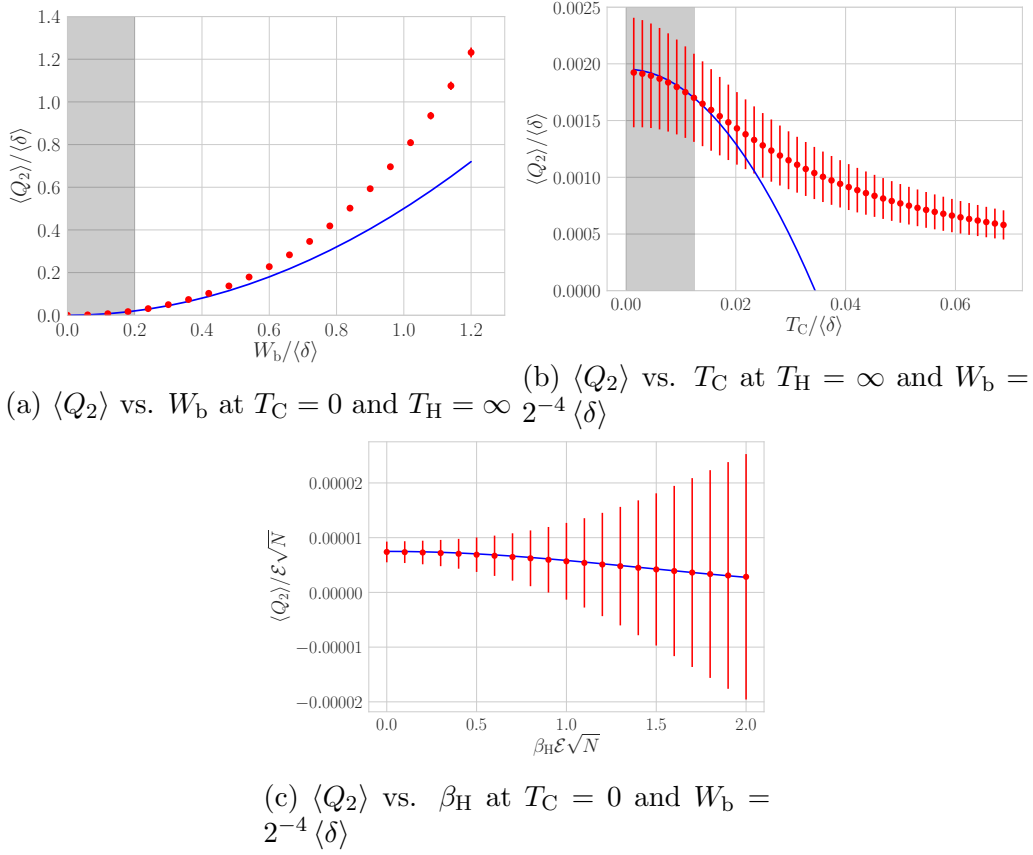


Figure 6.3: **Magnitude $|\langle Q_2 \rangle|$ of the average heat absorbed during cold thermalization (stroke 2) as a function of the cold-bath bandwidth W_b (6.3a), the cold-bath temperature T_C (6.3b), and the hot-bath temperature $T_H = 1/\beta_H$ (6.3c):** The blue lines represent the magnitude of the analytical prediction (6.18); in the bottom figure (showing β_H dependence) we use the full analytical result from [238]. See Sec. 6.7 for other parameters and definitions. The analytics match the numerics' shapes, and the agreement is fairly close, in the appropriate limits (where $\frac{W_b}{\langle \delta \rangle} \ll 1$ and $T_C / \langle \delta \rangle \ll 1$, in the gray shaded regions). The analytics systematically underestimate $\langle Q_2 \rangle$ at fixed W_b , due to the small level repulsion at finite N . The analytical prediction (6.18) substantially underestimates $\langle Q_2 \rangle$ when the cold-bath bandwidth is large, $W_b \gtrsim \langle \delta \rangle$. Such disagreement is expected: The analytics rely on $\frac{W_b}{\langle \delta \rangle} \ll 1$, neglecting chains of small gaps: $\delta'_j, \delta'_{j+1}, \dots < W_b$. Such chains proliferate as W_b grows. A similar reason accounts for the curve's crossing the origin in Fig. 6.3b: We analytically compute $\langle Q_2 \rangle$ only to second order in $T_C / \langle \delta \rangle$.

$Q_4 > 0$ to return to its start-of-trial state. Now, suppose that the cold bath is only cool: $T_C \gtrsim 0$. Cold thermalization might leave S in the upper level. S needs less heat, on average, to reset than if $T_C = 0$. A finite T_C therefore detracts from $\langle Q_4 \rangle$. The $+4 \ln 2 \frac{W_b}{\beta_C \langle \delta \rangle}$ offsets the detracting. However, the positive correction is smaller than the negative correction, as $\frac{W_b}{\langle \delta \rangle} \ll 1$.

A similar argument concerns $T_H < \infty$. But the β_H correction is too small to include in Eq. (6.21): $\langle Q_4 \rangle \approx W_b - \frac{2 \ln 2}{\beta_C} + \frac{(W_b)^2}{2 \langle \delta \rangle} e^{-N(\beta_H \mathcal{E})^2/4}$.

Figure 6.4 shows Eq. (6.21), to lowest order in T_C , as well as the β_H dependence of $\langle Q_4 \rangle$. The analytical prediction is compared with numerical simulations. The agreement is close, up to finite-size effects, in the appropriate regime ($T_C \ll W_b \ll \langle \delta \rangle$).

6.5.4 Average per-cycle power $\langle W_{\text{tot}} \rangle$

By the first law of thermodynamics, the net work outputted by the engine equals the net heat absorbed. Summing Eqs. (6.21) and (6.20) yields the per-trial power, or average work outputted per engine cycle:

$$\boxed{\langle W_{\text{tot}} \rangle} = \langle Q_2 \rangle + \langle Q_4 \rangle \approx W_b - \frac{2 \ln 2}{\beta_C} + 4 \ln 2 \frac{W_b}{\beta_C \langle \delta \rangle}. \quad (6.22)$$

The leading-order β_H correction is negative and too small to include—of order $\langle \delta \rangle \left(\frac{W_b}{\langle \delta \rangle} \right)^2 N (\beta_H \mathcal{E})^2$. Equation (6.22) agrees well with the numerics in the appropriate limits ($T_C \ll W_b \ll \langle \delta \rangle$) and beyond, as shown in Fig. 6.5. The main text contains the primary analysis of Eq. (6.22). Here, we discuss the $\langle Q_2 \rangle$ correction, limiting behaviors, and scaling.

The negative $\langle Q_2 \rangle = -\frac{(W_b)^2}{\langle \delta \rangle}$ detracts little from the leading term W_b of $\langle Q_4 \rangle$: $\frac{(W_b)^2}{\langle \delta \rangle} \ll W_b$, since $\frac{W_b}{\langle \delta \rangle} \ll 1$. The $\langle Q_2 \rangle$ cuts down on the per-trial power little.

The limiting behavior of Eq. (6.22) makes sense: Consider the limit as $W_b \rightarrow 0$. The cold bath has too small a bandwidth to thermalize the engine, so the engine should output no work, on average. Indeed, the first and third terms in Eq. (6.22) vanish, being proportional to W_b . The second term vanishes because $\beta_C \rightarrow \infty$ more quickly than $W_b \rightarrow 0$: The cold bath is very cold.

Equation (6.22) scales with the system size N no more quickly than $\sqrt{N}/2^N$, by the assumption $W_b \ll \langle \delta \rangle \sim \sqrt{N}/2^N$. This scaling makes sense: The engine outputs work because the energy eigenvalues meander upward and downward in Fig. 6.2 as $H(t)$ is tuned. In the thermodynamic limit, levels squeeze to-

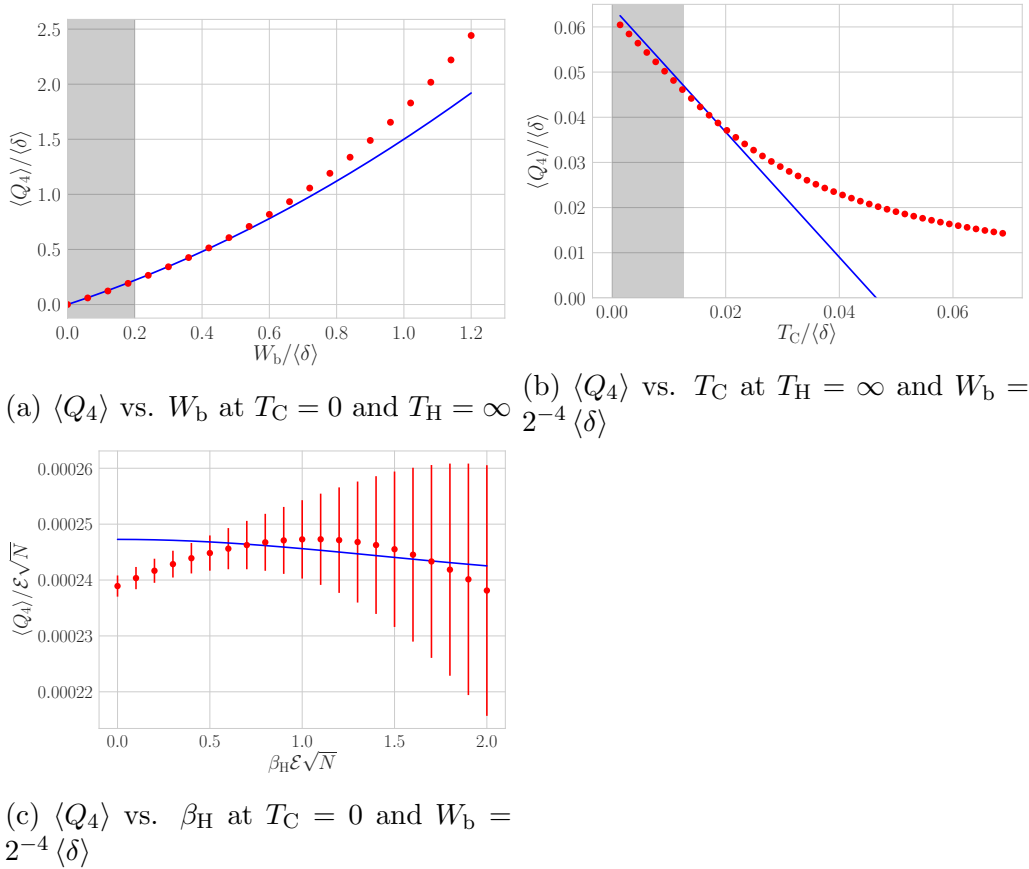


Figure 6.4: **Average heat $\langle Q_4 \rangle$ absorbed during hot thermalization (stroke 4) as a function of the cold-bath bandwidth W_b , the cold-bath temperature T_C , and the hot-bath temperature $T_H = 1/\beta_H$:** The blue lines represent the analytical prediction (6.21), to lowest order in T_C , with the β_H dependence of $\langle Q_4 \rangle$, too small a correction to include in Eq. (6.21): $\langle Q_4 \rangle \approx W_b - \frac{2 \ln 2}{\beta_C} + \frac{(W_b)^2}{2 \langle \delta \rangle} e^{-N(\beta_H \varepsilon)^2/4}$. In the bottom figure (showing β_H dependence) we again use the full analytical result from [238]. See Sec. 6.7 for other parameters and definitions. The analytics' shapes agree with the numerics', and the fit is fairly close, in the appropriate limits (where $e^{-\beta_C W_b} \ll 1$, $\frac{1}{\beta_C \langle \delta \rangle} \ll 1$, and $\frac{W_b}{\langle \delta \rangle} \ll 1$, in the gray shaded regions). The predictions underestimate $\langle Q_4 \rangle$; see the Fig. 6.3 caption. Figure 6.4c suggests that the numerics deviate significantly from the analytics: The numerics appear to depend on β_H via a linear term absent from the $\langle Q_4 \rangle$ prediction. This seeming mismatch appears symptomatic of finite sample and system sizes.


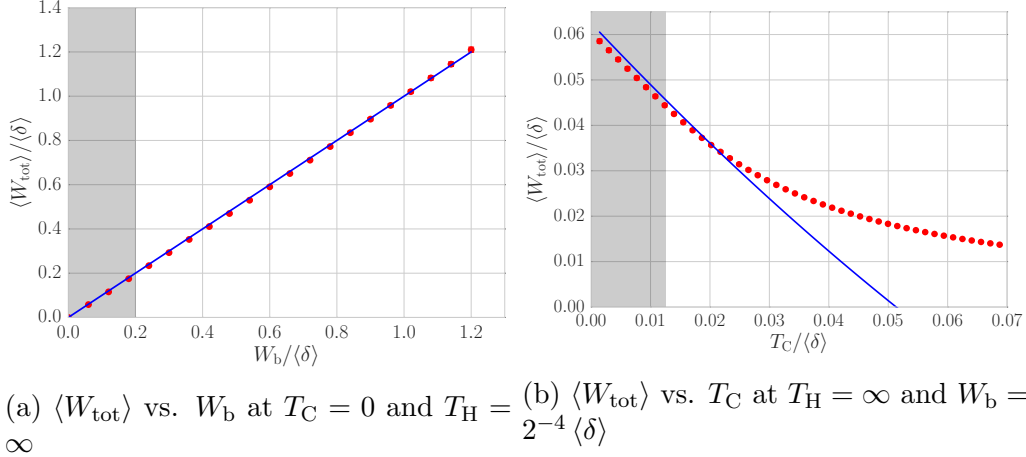


Figure 6.5: **Per-cycle power $\langle W_{\text{tot}} \rangle$ as a function of the cold-bath bandwidth W_b , the cold-bath temperature T_C , and the hot-bath temperature $T_H = 1/\beta_H$:** The blue lines represent the analytical prediction $\langle W_{\text{tot}} \rangle \approx W_b - \frac{2 \ln 2}{\beta_C}$: Eq. (6.22), to first order in $\frac{W_b}{\langle \delta \rangle}$ and in $\frac{1}{\beta_C \langle \delta \rangle}$. See Sec. 6.7 for other parameters and definitions. The analytics largely agree with the numerics in the appropriate regime: $\frac{W_b}{\langle \delta \rangle} \ll 1$, and $\frac{T_C}{\langle \delta \rangle} \ll 1$ (in the gray shaded region). Outside that regime, the analytics underestimate $\langle W_{\text{tot}} \rangle$; see Fig. 6.3 for an analysis. Figure 6.5c suggests that the numerics depend on β_H via a linear term absent from the analytical prediction; see the caption of Fig. 6.4c.

gether. Energy eigenvalues have little room in which to wander, and S outputs little work—hence our parallelization of fixed-length mesoscopic subengines in the thermodynamic limit (Sec. 6.6).

6.5.5 Efficiency η_{MBL} in the adiabatic approximation

The efficiency is defined as

$$\eta_{\text{MBL}} := \frac{\langle W_{\text{tot}} \rangle}{\langle Q_{\text{in}} \rangle}. \quad (6.23)$$

The numerator is averaged separately from the denominator because averaging W_{tot} over runs of one mesoscopic engine is roughly equivalent to averaging over simultaneous runs of parallel subengines in one macroscopic engine. $\frac{\langle W_{\text{tot}} \rangle}{\langle Q_{\text{in}} \rangle}$ may therefore be regarded as the $\frac{W_{\text{tot}}}{Q_{\text{in}}}$ of one macroscopic-engine trial.

The positive-heat-absorbing-stroke is stroke 4, in the average trial:

$$\langle Q_{\text{in}} \rangle = \langle Q_4 \rangle = \langle W_{\text{tot}} \rangle - \langle Q_2 \rangle = \langle W_{\text{tot}} \rangle \left(1 - \frac{\langle Q_2 \rangle}{\langle W_{\text{tot}} \rangle} \right) = \langle W_{\text{tot}} \rangle (1 + \phi) , \quad (6.24)$$

wherein

$$\phi := -\frac{\langle Q_2 \rangle}{\langle W_{\text{tot}} \rangle} \approx \frac{W_{\text{b}}}{2 \langle \delta \rangle} . \quad (6.25)$$

Substituting from Eq. (6.24) into Eq. (6.23) yields

$$\boxed{\eta_{\text{MBL}} \approx} \frac{\langle W_{\text{tot}} \rangle}{\langle W_{\text{tot}} \rangle (1 + \phi)} \approx 1 - \phi = \boxed{1 - \frac{W_{\text{b}}}{2 \langle \delta \rangle}} . \quad (6.26)$$

Using suboptimal baths diminishes the efficiency. Adding β_{C} -dependent terms from Eq. (6.22) to $\langle W_{\text{tot}} \rangle$ yields

$$\phi' = \frac{W_{\text{b}}}{2 \langle \delta \rangle} + \frac{\ln 2}{\beta_{\text{C}} \langle \delta \rangle} - 2 \ln 2 \frac{W_{\text{b}}}{\langle \delta \rangle} \frac{1}{\beta_{\text{C}} \langle \delta \rangle} . \quad (6.27)$$

The β_{H} correction, $1 - \frac{W_{\text{b}}}{2 \langle \delta \rangle} e^{-N(\beta_{\text{H}} \mathcal{E})^2/4}$, is too small to include. The correction shares the sign of β_{H} : A lukewarm hot bath lowers the efficiency.

Expressions (6.26) and (6.27) are compared with results from numerical simulations in Fig. 6.6. The analytics agree with the numerics in the appropriate regime ($T_{\text{C}} \ll W_{\text{b}} \ll \langle \delta \rangle$).

6.6 MBL engine in the thermodynamic limit

In Secs. 6.2, 6.5, and 6.4 we analyzed a single mesoscale engine. The power output of that mesoscale engine was very small, and indeed decreased exponentially with the size of the engine: $\langle W_{\text{tot}} \rangle \sim W_{\text{b}} \ll \langle \delta \rangle \sim 2^{-N}$. Worse, the tuning speed v must shrink exponentially: $H_{\text{meso}}(t)$ is ideally tuned quantum-adiabatically. The time per tuning stroke must far exceed $\langle \delta \rangle^{-1}$.

The mesoscale engine scales poorly, but properties of MBL offer a solution.

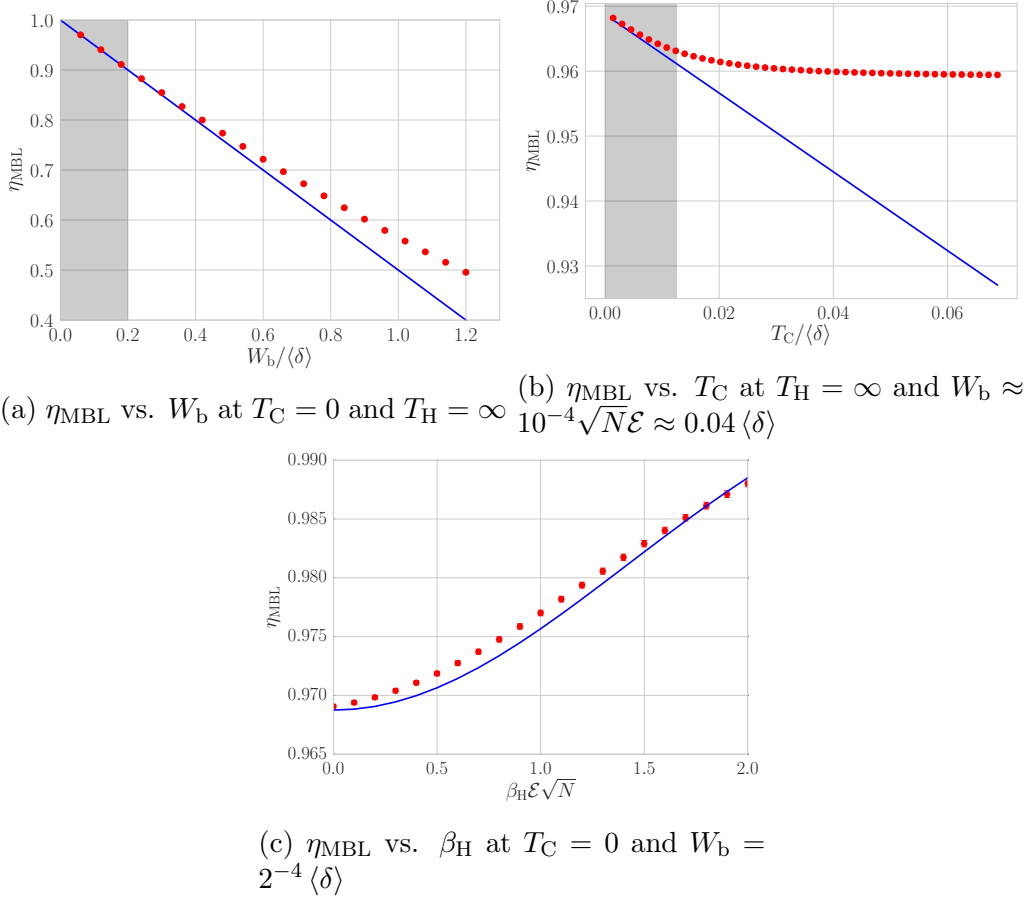


Figure 6.6: **Efficiency η_{MBL} as a function of the cold-bath bandwidth W_b , the cold-bath temperature T_C , and the hot-bath temperature $T_H = 1/\beta_H$:** The blue lines represent the analytical predictions (6.26) and (6.27). Figure (6.6c) shows the leading-order β_H dependence of η_{MBL} , a correction too small to include in Eq. (6.27): $1 - \frac{W_b}{2\langle\delta\rangle} e^{-N(\beta_H\mathcal{E})^2/4}$. See Sec. 6.7 for other parameters and definitions. The analytics agree with the numerics fairly well in the appropriate regime ($\frac{W_b}{\langle\delta\rangle} \ll 1$, $\frac{T_C}{\langle\delta\rangle} \ll 1$, and $\sqrt{N} T_H\mathcal{E} \ll 1$). The analytics underestimate η_{MBL} ; see the Fig. 6.3 caption.

The MBL engine’s advantage lies in having a simple thermodynamic limit that does not compromise efficiency or power output.

A thermodynamically large MBL Otto engine consists of mesoscale subengines that operate mostly independently. This independence hinges on *local level correlations* of the MBL phase [96, 196, 201]; cf Sec.1.1.4: Subsystems separated by a distance L evolve roughly independently until times exponential in L , due to the localization [158].

Consider perturbing an MBL system locally. In the Heisenberg picture, the perturbing operator spreads across a distance logarithmic in time [158]. (See also [109].) To see this, recall from Sec. 1.1.3 that the gaps between states differing on L sites are roughly

$$\delta \sim \mathcal{E} 2^{-L} e^{-L/\xi} . \quad (6.28)$$

On a timescale t the system is indifferent to gaps less than t^{-1} , for

$$(\mathcal{E}t)^{-1} \sim 2^{-L} e^{-L/\xi} \quad (6.29)$$

or

$$L(t) \sim (\ln 2 + \xi^{-1})^{-1} \ln(HScalet) . \quad (6.30)$$

The longer the time t for which the perturbation lasts, the farther the influence spreads.

First imagine tuning an MBL Hamiltonian infinitely slowly, to preclude diabatic transitions: $t \rightarrow \infty$. Even if the Hamiltonian consists of spatially local terms, the perturbation to each term spreads across the lattice. The global system cannot be subdivided into independent subengines.⁷ The global system’s average gap vanishes in the thermodynamic limit: $\langle \delta \rangle \rightarrow 0$. Since $\langle W_{\text{tot}} \rangle \sim W_b \ll \langle \delta \rangle$, the per-cycle power seems to vanish in the thermodynamic limit: $W_b \rightarrow 0$.

But now imagine tuning the Hamiltonian at a finite speed v . How big of a region can the tuning of a single local term affect? Dimensional analysis

⁷Granted, subengines are coupled together even if the Hamiltonian is quenched infinitely quickly: $H_{\text{sim}}(t)$ encodes a nearest-neighbor interaction, for example. That interaction might be regarded as coupling the edge of subengine k with the edge of subengine $k + 1$. But subengines’ edges may be regarded as ill-defined. The sites definitively in subengine k , near subengine k ’s center, should not couple to the sites near subengine ℓ ’s center, for any $\ell \neq k$, if the subengines are to function mostly independently. Additionally, one may separate subengines with “fallow” buffer zones, as we do.

suggests that the relevant time scale is $t \sim \frac{\mathcal{E}}{v}$. Local perturbations then affect a region of length $\sim L(\mathcal{E}/v) \sim \zeta \ln(\mathcal{E}^2/v)$. On a length scale $L(\mathcal{E}/v)$, global level correlations govern the engine's performance less than local level correlations do, i.e., less than $R(L(\mathcal{E}/v), \omega)$ does. This correlator registers level repulsion at a scale independent of N . Finite-speed tuning renders finite the average gap accessible to independent subengines, the $\langle \delta \rangle$ that would otherwise close in the thermodynamic limit. Each mesoscale subengine therefore outputs $\langle W_{\text{tot}} \rangle > 0$.

We can explain the gap's finiteness differently: Suppose that the engine's state starts some trial with weight on the j th energy level. The eigenenergies wiggle up and down during stroke 1. The j th energy may approach the $(j - 1)$ th. Such close-together energies likely correspond to far-apart subengines. If the levels narrowly avoided crossing, particles would be rearranged across a large distance. Particles must not be, as subengines must function independently. So the engine must undergo a diabatic transition: the engine's state must retain its configuration. The engine must behave as though the approaching energy level did not exist. Effectively removing the approaching level from the available spectrum effectively creates a gap in the spectrum. One can create such an effective gap (can promote such diabatic transitions) by tuning the Hamiltonian at a finite v .

Let us apply this principle to a chain of N -site mesoscale engines separated by N -site buffers. The engine is cycled between a shallowly localized (H_{GOE} -like) Hamiltonian, which has a localization length $\xi_{>}$, and a deeply localized (H_{MBL} -like) Hamiltonian, which has $\xi_{<} \ll \xi_{>}$.

The key element in the construction is that the cold bath acts through local operators confined to $< N \sim \xi_{>}$ sites. This defines the subengines of the thermodynamic MBL Otto engine. Localization guarantees that "what happens in a subengine stays in a subengine": subengines do not interfere much with each other's operation.

This subdivision boosts the engine's power. A length- N mesoscale engine operates at the average per-cycle power $\langle W_{\text{tot}} \rangle_{\text{meso}} \sim W_{\text{b}} \ll \frac{\mathcal{E}\sqrt{N}}{2^N}$ (Sec. 6.5). A subdivided length- N_{macro} MBL engine outputs average work $\sim \frac{N_{\text{macro}}}{2^N} \langle W_{\text{tot}} \rangle_{\text{meso}}$. In contrast, if the length- N_{macro} engine were not subdivided, it would output average work $\sim \frac{\mathcal{E}\sqrt{N_{\text{macro}}}}{2^{N_{\text{macro}}}}$, which vanishes in the thermodynamic limit.

We can be more quantitative.

6.7 Numerical simulations

We use numerical exact diagonalization to check analytical results. In Sec. 6.7.1, we describe the Hamiltonian used in our numerics. In Sec. 6.7.6, we study engine performance in the adiabatic limit (addressed analytically in Sec. 6.5). The code is available at <https://github.com/christopherdavidwhite/MBL-mobile>.

Call the times at which the strokes end $t = \tau, \tau', \tau'',$ and τ''' (see Fig. 6.1). For each of $N_{\text{reals}} \approx 1,000$ disorder realizations, we computed the whole density matrix $\rho(t)$ at $t = 0, \tau, \tau', \tau'', \tau'''$. The engine's time- t internal energy is $E(t) = \text{Tr}(H(t)\rho(t))$. The quantities of interest are straightforwardly

$$\langle W_1 \rangle = E(0) - E(\tau), \quad \langle W_3 \rangle = E(\tau''') - E(\tau''), \quad (6.31)$$

$$\langle Q_2 \rangle = E(\tau'') - E(\tau'), \quad \text{and} \quad \langle Q_4 \rangle = E(0) - E(\tau'''). \quad (6.32)$$

We disorder-averaged these quantities before dividing to compute the efficiency, $\eta_{\text{MBL}} = 1 - \frac{\langle W_1 \rangle + \langle W_3 \rangle}{\langle Q_4 \rangle}$.

6.7.1 Hamiltonian

The engine can be implemented with a disordered Heisenberg model. A similar model's MBL phase has been realized with cold atoms [188]. We numerically simulated a 1D mesoscale chain governed by a Hamiltonian

$$H_{\text{sim}}(t) = \frac{\mathcal{E}}{Q(h(\alpha_t))} \left[\sum_{j=1}^{N-1} \vec{\sigma}_j \cdot \vec{\sigma}_{j+1} + h(\alpha_t) \sum_{j=1}^N h_j \sigma_j^z \right]; \quad (6.33)$$

this is a special case of the general mesoscopic Hamiltonian (6.2) described in Sec. 6.2. Equation (6.33) describes spins equivalent to interacting spinless fermions. Energies are expressed in units of \mathcal{E} , the average per-site energy density. For $\gamma = x, y, z$, the γ th Pauli operator that operates nontrivially on the j th site is denoted by σ_j^γ . The Heisenberg interaction $\vec{\sigma}_j \cdot \vec{\sigma}_{j+1}$ encodes nearest-neighbor hopping and repulsion.

The tuning parameter $\alpha_t \in [0, 1]$ determines the phase occupied by $H_{\text{sim}}(t)$. The site- j disorder potential depends on a random variable h_j distributed uniformly across $[-1, 1]$. The disorder strength $h(\alpha_t)$ varies as $h(\alpha_t) = \alpha_t h_{\text{GOE}} + (1 - \alpha_t)h_{\text{MBL}}$. When $\alpha_t = 0$, the disorder is weak, $h = h_{\text{GOE}}$, and the engine occupies the ETH phase. When $\alpha_t = 1$, the disorder is strong, $h = h_{\text{MBL}} \gg h_{\text{GOE}}$, and the engine occupies the MBL phase.

The normalization factor $Q(h(\alpha_t))$ preserves the width of the density of states (DoS) and so preserves $\langle \delta \rangle$. $Q(h(\alpha_t))$ prevents the work extractable via change of bandwidth from polluting the work extracted with help from level statistics (see Sec. 8.1.1 for a discussion of work extraction from bandwidth change). $Q(h(\alpha_t))$ is defined and calculated below.

The ETH-side field had a magnitude $h(0) = 2.0$, and the MBL-side field had a magnitude $h(1) = 20.0$. These $h(\alpha_t)$ values fall squarely on opposite sides of the MBL transition at $h \approx 7$.

6.7.2 Scaling factor

We wish to keep the DoS constant through the cycle. To fix $\mu(E)$, we rescale the Hamiltonian by a factor $Q(h(\alpha_t))$. We define $Q^2(h(\alpha_t))$ as the disorder average of the variance of the unrescaled DoS:

$$Q^2(h(\alpha_t)) := \left\langle \left(\frac{1}{\mathcal{N}} \sum_{j=1}^{\mathcal{N}} E_j^2 \right) - \left(\frac{1}{\mathcal{N}} \sum_{j=1}^{\mathcal{N}} E_j \right)^2 \right\rangle_{\text{disorder}} \quad (6.34)$$

$$= \left\langle \frac{1}{\mathcal{N}} \text{Tr}(\tilde{H}^2(t)) - \left(\frac{1}{\mathcal{N}} \text{Tr}(\tilde{H}(t)) \right)^2 \right\rangle_{\text{disorder}}. \quad (6.35)$$

The $\tilde{H}(t)$ denotes an unrescaled variation on the random-field Heisenberg Hamiltonian $H(t)$ of Eq. (6.33):

$$\tilde{H}(t) := \mathcal{E} \left[\sum_{j=1}^{N-1} \vec{\sigma}_j \cdot \vec{\sigma}_{j+1} + h(\alpha_t) \sum_{j=1}^N h_j \sigma_j^z \right]. \quad (6.36)$$

To compute $Q^2(h(\alpha_t))$, we rewrite the unrescaled Hamiltonian as

$$\tilde{H}(t) = \mathcal{E} \left[2 \sum_{j=1}^{N-1} (\sigma_j^+ \sigma_{j+1}^- + \text{h.c.}) + \sum_{j=1}^{N-1} \sigma_j^z \sigma_{j+1}^z + h(\alpha_t) \sum_{j=1}^N h_j \sigma_j^z \right]. \quad (6.37)$$

We assume that N is even, and we work at half-filling. The $\frac{N}{2}$ -particle subspace has dimensionality $\mathcal{N} = \binom{N}{N/2}$.

Let us calculate some operator traces that we will invoke later. Let $X := \prod_{j=1}^N \sigma_j^x$ denote the global spin-flip operator. For any operator A such that $X^\dagger A X = -A$,

$$\text{Tr}(A) = \text{Tr}(X^\dagger A X) = -\text{Tr}(A). \quad (6.38)$$

We have used the evenness of N , which implies the invariance of the half-filling subspace under X . Also, $\text{Tr}(A) = 0$. In particular, $0 = \text{Tr}(\sigma_j^z) = \text{Tr}(\sigma_j^z \sigma_{j'}^z \sigma_{j''}^z)$, if $j \neq j' \neq j''$.

Traces of products of even numbers of σ^z factors require more thought:

$$\text{Tr}(\sigma_j^z \sigma_{j+1}^z) = (\# \text{ states } j, j+1 = \uparrow\uparrow) + (\# \text{ states } j, j+1 = \downarrow\downarrow) \quad (6.39)$$

$$- 2(\# \text{ states } j, j+1 = \uparrow\downarrow)$$

$$= \binom{N-2}{N/2-2} + \binom{N-2}{N/2} - 2 \binom{N-2}{N/2-1}$$

$$= -\mathcal{N} \frac{1}{N-1}. \quad (6.40)$$

Similarly,

$$\text{Tr}([\sigma_j^+ \sigma_j^-][\sigma_{j+1}^- \sigma_{j+1}^+]) = \text{Tr}([\sigma_j^- \sigma_j^+][\sigma_{j+1}^+ \sigma_{j+1}^-])$$

$$= (\# \text{ states } j, j+1 = \uparrow\downarrow)$$

$$= \binom{N-2}{N/2-1} \quad (6.41)$$

$$= \mathcal{N} \frac{N}{4(L-1)}, \quad (6.42)$$

and

$$\text{Tr}(\sigma_j^z \sigma_{j+1}^z \sigma_{j'}^z \sigma_{j'+1}^z) = (\# \text{ states } j, j+1, j', j'+1 = \uparrow\uparrow\uparrow\uparrow)$$

$$+ \binom{4}{2} (\# \text{ states } j, j+1, j', j'+1 = \uparrow\uparrow\downarrow\downarrow)$$

$$+ (\# \text{ states } j, j+1, j', j'+1 = \downarrow\downarrow\downarrow\downarrow)$$

$$- \binom{4}{1} (\# \text{ states } j, j+1, j', j'+1 = \uparrow\uparrow\uparrow\downarrow)$$

$$- \binom{4}{1} (\# \text{ states } j, j+1, j', j'+1 = \uparrow\downarrow\downarrow\downarrow)$$

$$= \binom{N-4}{N/2-4} + 6 \binom{N-4}{N/2-2} + \binom{N-4}{N/2} \quad (6.43)$$

$$- 6 \binom{N-4}{N/2-3} - 6 \binom{N-4}{N/2-1}$$

$$= \mathcal{N} \frac{3}{(N-1)(N-3)}, \quad (6.44)$$

wherein the first equality's combinatorial factors come from permutations on

sites j , $j + 1$, j' , and $j' + 1$.

Assembling these pieces, we find $\text{Tr}(\tilde{H}(t)) = \mathcal{E} \sum_{j=1}^{N-1} \text{Tr}(\sigma_j^z \sigma_j^z) = -\mathcal{E}\mathcal{N}$. Next, we compute $\text{Tr}(\tilde{H}^2(t))$:

$$\begin{aligned}
\tilde{H}^2(t) &= \mathcal{E}^2 \left[4 \sum_j^{N-1} (\sigma_j^+ \sigma_j^-) (\sigma_{j+1}^- \sigma_{j+1}^+) + 4 \sum_j^{N-1} (\sigma_j^- \sigma_j^+) (\sigma_{j+1}^+ \sigma_{j+1}^-) \right. \\
&\quad \left. + \sum_{j,j'=1}^{N-1} \sigma_j^z \sigma_{j+1}^z \sigma_{j'}^z \sigma_{j'+1}^z + h^2(\alpha_t) \sum_{j=1}^N h_j^2 + (\text{traceless terms}) \right] \quad (6.45) \\
&= \mathcal{E}^2 \left[4 \sum_j^{N-1} (\sigma_j^+ \sigma_j^-) (\sigma_{j+1}^- \sigma_{j+1}^+) + 4 \sum_j^{N-1} (\sigma_j^- \sigma_j^+) (\sigma_{j+1}^+ \sigma_{j+1}^-) \right. \\
&\quad + \sum_{j=1}^{N-1} I + \sum_{j=1}^{N-2} \sigma_j^z \sigma_{j+2}^z + \sum_{j=1}^{N-3} \sum_{j'=j+2}^{N-1} \sigma_j^z \sigma_{j+1}^z \sigma_{j'}^z \sigma_{j'+1}^z \\
&\quad \left. + h(\alpha_t)^2 (\alpha_t) \sum_{j=1}^N h_j^2 + (\text{traceless terms}) \right]. \quad (6.46)
\end{aligned}$$

We take the trace, using Eqs. (6.39), (6.41), and (6.43):

$$\text{Tr}(\tilde{H}^2(t)) = \mathcal{N} \left[3N - 1 + \frac{N-2}{N-1} + h^2 \sum_{j=1}^N h_j^2 \right]. \quad (6.47)$$

We disorder-average by taking $h_j^2 \mapsto \int_0^1 dh_j h_j^2 = \frac{1}{3}$:

$$\left\langle \text{Tr}(\tilde{H}^2(t)) \right\rangle_{\text{disorder}} = \mathcal{N} \left[3N - 1 + \frac{N-2}{N-1} + N \frac{h^2}{3} \right]. \quad (6.48)$$

Substituting into Eq. (6.34), we infer the rescaling factor's square:

$$Q^2(h(\alpha_t)) = 3N - 2 + \frac{N-2}{N-1} + N \frac{h^2}{3}. \quad (6.49)$$

Our results are insensitive to the details of Q . The width of the DoS in one disorder realization will differ from the disorder average (6.49). Moreover, that difference will vary as we tune $h(\alpha_t)$, because the disorder affects only one term. The agreement between the analytics, in which $\mu(E)$ is assumed to remain constant in t , and the numerics is therefore comforting: the engine is robust against small variations in the rescaling.

6.7.3 Representing states and Hamiltonians

We structured our software to facilitate a possible extension: The cold bath might be modeled more realistically, as coupling to the engine only locally.

We represent the state of one mesoscopic MBL Otto engine with a density matrix $\rho \in \mathbb{C}^{\mathcal{N} \times \mathcal{N}}$, and the Hamiltonian with a matrix $H \in \mathbb{C}^{\mathcal{N} \times \mathcal{N}}$, relative to the basis $\{|s_1\rangle, \dots, |s_{\mathcal{N}}\rangle\} = \{|\uparrow \dots \uparrow\rangle, \dots, |\downarrow \dots \downarrow\rangle\}$ of products of σ^z eigenstates. We track the whole density matrix, rather than just the energy-diagonal elements, with an eye toward the coherent superpositions that diabatic corrections create. For an N -site chain at half-filling, $\mathcal{N} = \binom{N}{N/2} \simeq \sqrt{\frac{2}{\pi N}} 2^N$.

6.7.4 Strokes 1 and 3: tuning

The (l, m) entry of the initial-state density matrix is

$$\rho(0)_{lm} = \langle s_l | \frac{1}{Z} e^{-\beta_{\text{H}} H(0)} | s_m \rangle = \frac{1}{Z} \sum_j e^{-\beta_{\text{H}} E_j(0)} \langle s_l | E_j(0) \rangle \langle E_j(0) | s_m \rangle. \quad (6.50)$$

The j th eigenstate of $H(0)$, associated with energy $E_j(0)$, is denoted by $|E_j(0)\rangle$. We approximate the time evolution from 0 to τ (during stroke 1) as adiabatic. The evolution therefore does not move weight between levels:

$$\rho(\tau)_{lm} = \frac{1}{Z} \sum_j e^{-\beta_{\text{H}} E_j(0)} \langle s_l | E_j(\tau) \rangle \langle E_j(\tau) | s_m \rangle. \quad (6.51)$$

If we represented our density matrix relative to an instantaneous energy eigenbasis, simulating the time evolution would be trivial: We would reinterpret the diagonal matrix ρ as being diagonal, with the same elements in a new basis. However, we wish to represent $\rho(t)$ relative to the σ_j^z product basis. This representation enhances the code's flexibility, facilitating the inclusion of diabatic evolutions and a more detailed model of cold thermalization. To represent $\rho(t)$ relative to the σ_j^z product basis, we note that

$$\rho(\tau)_{lm} = \sum_j \langle s_l | E_j(\tau) \rangle \langle E_j(0) | \rho(0) | E_j(0) \rangle \langle E_j(\tau) | s_m \rangle = [U(\tau, 0) \rho(0) U(\tau, 0)^\dagger]_{lm}. \quad (6.52)$$

We have defined a time-evolution matrix $U(\tau, 0) \in \mathbb{C}^{\mathcal{N} \times \mathcal{N}}$ by $U(\tau, 0)_{lm} = \sum_j \langle s_l | E_j(\tau) \rangle \langle E_j(0) | s_m \rangle$. This matrix is easily computed via exact diagonalization of $H(0)$ and $H(\tau)$.

We can compute the density matrix $\rho(\tau'')$ at the end of stroke 3 (the tuning from MBL to GOE) from the density matrix $\rho(\tau')$ at the end of stroke 2

(the cold-bath thermalization) similarly: $\rho(\tau'') = U(\tau'', \tau')\rho(\tau')U(\tau'', \tau')^\dagger$. The time-evolution matrix $U(\tau'', \tau') \in \mathbf{C}^{\mathcal{N} \times \mathcal{N}}$ is given by

$$U(\tau'', \tau')_{lm} = \sum_j \langle s_l | E_j(0) \langle E_j(\tau) \rangle s_m . \quad (6.53)$$

[Recall that $H(\tau'') = H(0)$ and $H(\tau') = H(\tau)$.]

6.7.5 Stroke 2: Thermalization with the cold bath

During stroke 2, the system thermalizes with a bandwidth- W_b cold bath. We make three assumptions. First, the bandwidth cutoff is hard: The bath can transfer only amounts $< W_b$ of energy at a time. Therefore, the cold bath cannot move probability mass between adjacent levels separated by just one gap $\delta' > W_b$. Second, the bath is Markovian. Third, the system thermalizes for a long time. The bath has time to move weight across sequences of small gaps $\delta'_j, \delta'_{j+1}, \dots < W_b$.

We can implement thermalization as follows. First, we identify sequences of levels connected by small gaps. Second, we reapportion weight amongst the levels according to a Gibbs distribution.

Suppose, for example, that the MBL Hamiltonian $H(\tau)$ contains the following chain of six energies, E_1, \dots, E_6 , separated from its surrounding levels by large gaps (Fig. 6.7):

$$\begin{aligned} (E_2 - E_1), (E_3 - E_2) &< W_b, \\ (E_5 - E_4) &< W_b, \quad \text{and} \\ (E_4 - E_3), (E_6 - E_5) &> W_b. \end{aligned} \quad (6.54)$$

We suppress the time arguments to simplify notation. Before thermalization, the density operator is diagonal with respect to the energy basis: $\rho(\tau) = \sum_j \rho_j |E_j\rangle\langle E_j|$. The weight on level j is denoted by ρ_j . Thermalization maps

$$\begin{aligned} \rho(\tau) \mapsto \rho(\tau') &= \frac{\rho_1 + \rho_2 + \rho_3}{e^{-\beta_C E_1} + e^{-\beta_C E_2} + e^{-\beta_C E_3}} \\ &\quad \times \left(e^{-\beta_C E_1} |E_1\rangle\langle E_1| + e^{-\beta_C E_2} |E_2\rangle\langle E_2| + e^{-\beta_C E_3} |E_3\rangle\langle E_3| \right) \\ &+ \frac{\rho_4 + \rho_5}{e^{-\beta_C E_4} + e^{-\beta_C E_5}} \left(e^{-\beta_C E_4} |E_4\rangle\langle E_4| + e^{-\beta_C E_5} |E_5\rangle\langle E_5| \right) \\ &+ \rho_6 |E_6\rangle\langle E_6|. \end{aligned} \quad (6.55)$$




Figure 6.7: **Energies of a cold-thermalized system:** We illustrate our implementation of cold thermalization with this example chain of six energies. The cold bath has a bandwidth of size W_b , depicted in green.

6.7.6 Results

We compare the analytical predictions of Sec. 6.5 to numerical simulations of a 12-site engine governed by the Hamiltonian (6.33). During strokes 1 and 3, the state was evolved as though the Hamiltonian were tuned adiabatically. We index the energies $E_j(t)$ from least to greatest at each instant: $E_j(t) < E_k(t) \forall j < k$. Let ρ_j denote the state's weight on eigenstate j of the initial Hamiltonian, whose $\alpha_t = 0$. The engine ends stroke 1 with weight ρ_j on eigenstate j of the post-tuning Hamiltonian, whose $\alpha_t = 1$.

The main results appear in Fig. 6.8. Figure 6.8a shows the average work extracted per cycle, $\langle W_{\text{tot}} \rangle$. Figure 6.8b shows the efficiency, η_{MBL} .

In these simulations, the baths had the extreme temperatures $T_H = \infty$ and $T_C = 0$. This limiting case elucidates the W_b -dependence of $\langle W_{\text{tot}} \rangle$ and of η_{MBL} : Disregarding finite-temperature corrections, on a first pass, builds intuition. Finite-temperature numerics appear alongside finite-temperature analytical calculations above.

Figure 6.8 shows how the per-cycle power and the efficiency depend on the cold-




Figure 6.8: **Average per-cycle power $\langle W_{\text{tot}} \rangle$ (left) and efficiency η_{MBL} (right) as functions of the cold-bath bandwidth W_b :** Each red dot represents an average over 1,000 disorder realizations of the random-field Heisenberg Hamiltonian (6.33). The slanted blue lines represent the analytical predictions (6.7) and (6.9) of Sec. 6.5. When $W_b \ll \langle \delta \rangle$ (in the gray shaded region), $\langle W_{\text{tot}} \rangle$ and η_{MBL} vary linearly with W_b , as predicted. The error bars are smaller than the numerical-data points.

bath bandwidth W_b . As expected, $\langle W_{\text{tot}} \rangle \approx W_b$. The dependence’s linearity, and the unit proportionality factor, agree with Eq. (6.7). Also as expected, the efficiency declines as the cold-bath bandwidth rises: $\eta_{\text{MBL}} \approx 1 - \frac{W_b}{2\langle \delta \rangle}$. The linear dependence and the proportionality factor agree with Eq. (6.9).

The gray columns in Fig. 6.8 highlight the regime in which the analytics were performed, where $\frac{W_b}{\langle \delta \rangle} \ll 1$. If the cold-bath bandwidth is small, $W_b < \langle \delta \rangle$, the analytics-numerics agreement is close. But the numerics agree with the analytics even outside this regime. If $W_b \gtrsim \langle \delta \rangle$, the analytics slightly underestimate η_{MBL} : The simulated engine operates more efficiently than predicted. To predict the numerics’ overachievement, one would calculate higher-order corrections in Sec. 6.5: One would Taylor-approximate to higher powers, modeling subleading physical processes. Such processes include the engine’s dropping across a chain of three small gaps, $\delta'_1, \delta'_2, \delta'_3 < W_b$, during cold thermalization.

The error bars are smaller than the numerical-data points. Each error bar represents the error in the estimate of a mean (of $\langle W_{\text{tot}} \rangle$ or of $\eta_{\text{MBL}} := 1 - \frac{\langle W_{\text{tot}} \rangle}{\langle Q_{\text{in}} \rangle}$) over 1,000 disorder realizations. Each error bar extends a distance (sample standard deviation)/ $\sqrt{\# \text{ realizations}}$ above and below that mean.

THE MBL-MOBILE AT FINITE SPEED

In Ch. 6, we assumed that we could tune the engine in strokes 2 and 4 infinitely slowly, as far as the individual mesoscale engines were concerned—but then we also assumed that we could tune it quickly enough that information does not propagate between those individual mesoscale engines.

Are these two assumptions broadly consistent? That is, does a regime exist that is both *slow* with respect to unwanted transitions in a single mesoscale engine and *fast* with respect to transitions across multiple mesoscale engines?

More generally, an experimenter would want to run our cycle in a finite time: she is restricted by the coherence times of her systems, by a desire to go to lunch, and ultimately by the requirement that she finish, publish, and graduate. How does our engine perform at finite speed?

The time scales of a macroscopic engine are crucial for the assessment of the MBL Otto engine. In this chapter we estimate the restrictions on the speed with which the Hamiltonian must be tuned to avoid diabatic transitions (Sec. 7.1). We compare these diabatic corrections to numerical simulations (Sec. 7.2) After that we compute the restrictions on the speed required to prevent intersubengine communication (Sec. 7.3); we find that we can indeed tune the system quickly enough to avoid communication, but slowly enough to avoid undue diabatic transitions. We finally estimate the time required for cold thermalization (stroke 2) in Sec. 7.4.

7.1 Diabatic corrections

We have modeled the Hamiltonian tuning as quantum-adiabatic, but realistic tuning speeds $v := \mathcal{E} \left| \frac{d\alpha_t}{dt} \right|$ are finite. To understand diabatic tuning's effects, we distinguish the time- t density matrix $\rho(t)$ from the corresponding diagonal ensemble,

$$\begin{aligned} \rho_{\text{diag}}(t) &= \sum_j |E_j(t)\rangle \varepsilon_j \langle E_j(t)|, \quad \text{wherein} \\ \varepsilon_j &= \langle E_j(t)| \rho |E_j(t)\rangle \end{aligned} \tag{7.1}$$




Figure 7.1: **Three (times two) classes of diabatic transitions:** Hops to arbitrary energy levels, modeled with general adiabatic perturbation theory (APT), plague the ETH regime. Landau-Zener transitions and fractional-Landau-Zener transitions plague the many-body-localized regime.

and $|E_j(t)\rangle$ is an instantaneous energy eigenbasis of $H_{\text{meso}}(t) = \sum_j |E_j(t)\rangle E_j(t) \langle E_j(t)|$. The average energy depends on $\rho(t)$ only through $\rho_{\text{diag}}(t)$. [More generally, the state’s off-diagonal elements dephase under the dynamics. $\rho_{\text{diag}}(t)$ is “slow” and captures most of the relevant physics [44].]

In the adiabatic limit, $\varepsilon_j(t) = \varepsilon_j(0)$. We seek to understand how this statement breaks down when the tuning proceeds at a finite speed v . It is useful to think of “infinite-temperature thermalization” in the sense of this diagonal ensemble: Fast tuning may push the diagonal-ensemble weights $\varepsilon_j(t)$ towards uniformity—even though the process is unitary and the entropy $S = -\rho(t) \ln \rho(t)$ remains constant—thanks to the off-diagonal elements.

The effects of diabatic tuning appear in three distinct regimes, which we label “fractional-Landau-Zener,” “Landau-Zener,” and “APT” (Fig. 7.1). We estimate the average per-cycle work costs $\langle W_{\text{diab}} \rangle$ of diabatic jumps, guided by the numerics in Sec. 6.7. We focus on $T_{\text{H}} = \infty$ and $T_{\text{C}} = 0$, for simplicity. Since $T_{\text{H}} = \infty$, diabatic hops cannot bring $\rho_{\text{diag}}(t)$ closer to $I/2^N$ —cannot change the average energy—during stroke 1. Hence we focus on stroke 3.

7.1.1 Fractional-Landau-Zener transitions

At the beginning of stroke 3, nonequilibrium effects could excite the system back across the small gap to energy level j . The transition would cost work and would prevent the trial from outputting $W_{\text{tot}} > 0$. We dub this excitation a fractional-Landau-Zener (frac-LZ) transition. It could be suppressed by a sufficiently slow drive [46]. The effects, and the resultant bound on v , are simple to derive

Let the gap start stroke 3 at size δ and grow to a size $\Delta > \delta$. Because the two energy levels begin close together, one cannot straightforwardly apply the Landau-Zener formula. One must use the fractional-Landau-Zener result of De Grandi and Polkovnikov [46],

$$p_{\text{frac-LZ}}(\delta) \approx \frac{v^2(\delta_-)^2}{16} \left(\frac{1}{\delta^6} + \frac{1}{\Delta^6} \right) \approx \frac{v^2(\delta_-)^2}{16\delta^6}. \quad (7.2)$$

δ_- denotes the MBL level-repulsion scale, the characteristic matrix element introduced by a perturbation between eigenstates of an unperturbed Hamiltonian. We suppose that energy-level pairs with $p_{\text{frac-LZ}} \lesssim 1$ are returned to the infinite-temperature state from which the cold bath disturbed them. These pairs do not contribute to $\langle W_{\text{tot}} \rangle$. Pairs that contribute have $p_{\text{frac-LZ}} \ll 1$, i.e.,

$$\delta \gtrsim (v\delta_-)^{1/3}. \quad (7.3)$$

If the rest of the stroke is adiabatic, the average work performed during the cycle is

$$\langle W_{\text{tot}} \rangle \sim \langle Q_4 \rangle - \langle Q_2 \rangle - (v\delta_-)^{1/3}, \quad (7.4)$$

which results immediately in the correction

$$\langle W_{\text{diab,frac-LZ}} \rangle \sim (v\delta_-)^{1/3}. \quad (7.5)$$

This correction is negligible at speeds low enough that

$$v \ll \frac{(W_{\text{b}})^3}{\delta_-}. \quad (7.6)$$

7.1.2 Landau-Zener transitions

While the system is localized, the disturbances induced by the tuning $\frac{dH(t)}{dt}$ can propagate only a short distance l_v . The tuning effectively reduces the mesoscale engine to an effective length- l_v subengine. To estimate l_v , we compare the minimum gap of a length- l_v subsystem to the speed v :

$$\mathcal{E}^{-l_v} e^{-l_v/\xi} \sim \sqrt{v}. \quad (7.7)$$

This minimum gap—the closest that two levels are likely to approach—is given by the smallest level-repulsion scale, δ_- . δ_- characterizes the deeply localized

system, whose $\xi = \xi_<$. Consequently,

$$l_v \sim \frac{\ln(\mathcal{E}^2/v)}{2\left(\ln 2 + \frac{1}{\xi_<}\right)}. \quad (7.8)$$

Suppose that $l_v \leq N$, and consider a length- l_v effective subengine. In the adiabatic limit, $\langle W_{\text{tot}} \rangle$ does not depend on the engine's size. ($\langle W_{\text{tot}} \rangle$ depends only on the bath bandwidth $W_b \ll \langle \delta \rangle$.) To estimate how a finite v changes $\langle W_{\text{tot}} \rangle$, we consider the gaps $\delta < W_b$ of the size- l_v subengine. We divide the gaps into two classes:

1. Gaps connected by flipping 1-bits on a region of diameter $l < l_v$. The tuning is adiabatic with respect to these gaps, so they result in work output.
2. Gaps connected by flipping 1-bits on a region of diameter $l = l_v$. The tuning is resonant with these gaps and so thermalizes them, in the sense of the diagonal ensemble: The tuning makes the instantaneous-energy-eigenvector weights ε_j uniform, on average.

Type-1 gaps form a v -independent $O(1)$ fraction θ of the length- l_v subengine's short-length-scale gaps.¹ Type-2 gaps therefore make up a fraction $1 - \theta$. Hence Landau-Zener physics leads to a v -independent $O(1)$ diabatic correction $(1 - \theta)W_b$ to $\langle W_{\text{tot}} \rangle$, provided that v is high enough that $l_v < N$.

7.1.3 APT transitions

When the system is in the ETH phase (or has correlation length $\xi \sim N$), typical minimum gaps (points of closest approach) are still given by the level-repulsion scale, which is now $\langle \delta \rangle$. Hence one expects the tuning to be adiabatic if

$$v \ll \langle \delta \rangle^2. \quad (7.9)$$

This criterion is at least as stringent (depending on the system size) as the requirement (7.6) that fractional Landau-Zener transitions occur rarely. Both

¹We can estimate θ crudely. For a given diameter- l_v subset, each gap connected by a diameter- $(l_v - 1)$ operator can be made into a diameter- l_v gap: One flips the last (l_v) th 1-bit, for $\theta \sim 1/2$. This estimate neglects several combinatorial matters. A more detailed analysis would account for the two different diameter- $(l_v - 1)$ regions of a given length- l_v subengine, gaps connected by 1-bit flips in the intersections of those subengines, the number of possible diameter- l_v subengines of an N -site system, etc.




Figure 7.2: **Average per-cycle work as a function of tuning speed:** We numerically simulated 995 disorder realizations of the random-field Heisenberg Hamiltonian (6.33) for a system of $N = 8$ sites (red dots). The results are compared to the analytical estimate (6.7) for the adiabatic work output (blue line) and an empirical straight-line fit $W_{\text{tot}} = W_0 - W_1(v\delta_-)^{1/3}/W_b$ (black line). Errors in the estimate of the mean, computed as $(\text{sample standard deviation})/\sqrt{(\# \text{ realizations})}$, lead to error bars smaller than the numerical-data points.

fractional Landau-Zener transitions and APT transitions bound the cycle time τ_{cycle} less stringently than thermalization with the cold bath; hence a more detailed analysis of APT transitions would be gratuitous. Such an analysis would rely on the general adiabatic perturbation theory of De Grandi and Polkovnikov [46]; hence the moniker “APT transitions.”

7.2 Numerical simulations

We simulate the engine as in Ch. 6 (in particular Sec. 6.7)—with the exception that now in the work strokes, strokes 1 and 3, $H_{\text{sim}}(t)$ was tuned at finite speed v . We numerically study the preclusion of communication between mesoscale subengines (addressed analytically in Sec. 7.3) only insofar as these results follow from diabatic corrections: Limitations on computational power restricted the system size to 12 sites.

Computational limitations restricted the engine to 8 sites. (These simulations quickly become slow to run because our upper bounds on v scale as powers of $\langle \delta \rangle \sim 2^{-N}$.)

7.2.1 Simulating finite-time tuning

We simulate a stepwise tuning, taking

$$\alpha_t = \frac{\delta t \lfloor t/\delta t \rfloor}{T}. \quad (7.10)$$

δt denotes a time-step size, and $T \propto \frac{h_{\text{MBL}} - h_{\text{GOE}}}{v}$ denotes the time for which one tuning stroke lasts. This protocol is more violent than the protocols treated analytically: v is assumed to remain finite in the diabatic analytics. In the numerics, we tune by sudden jumps (for reasons of numerical convenience). We work at $T_{\text{H}} = \infty$ and $T_{\text{C}} = 0$ —again, to capture the essential physics without the complication of finite-temperature corrections.

To implement this stepwise tuning, we take

$$\alpha(t) = \frac{\delta t \lfloor t/\delta t \rfloor}{T}, \quad (7.11)$$

where δt denotes a time-step size and $T \propto (h_{\text{MBL}} - h_{\text{GOE}})/v$ denotes the total tuning time, and we compute a time-evolution unitary for the whole stroke by chaining together the unitaries for each time step. For stroke 1,

$$U(\tau, 0; v, \delta t) = e^{-iH(\tau-\delta t)\delta t} e^{-iH(\tau-2\delta t)\delta t} \dots e^{-iH(0)\delta t}, \quad (7.12)$$

with the number of time steps set by the speed. We use the time step $\delta t = 0.405 \langle \delta \rangle$, but our results are not sensitive to time step's size.

In judging the engine's effectiveness at a finite v , we must estimate the level-repulsion scale δ_- . We do this by diagonalizing 10^6 disorder realizations at the relevant disorder width, $h = 20$, for $N = 8$ sites. A histogram of the gaps is plotted in Fig. 7.3. We then visually estimate the point at which the distribution turns over. Our results are not sensitive to this value.

Figure 7.2 shows the average work output, $\langle W_{\text{tot}} \rangle$, as a function of v . Despite the simulated protocol's violence, both a fractional-Landau-Zener correction $W_{\text{frac-LZ}} \sim (v\delta_-)^3$, explained in Sec. 7.1.1, and a v -independent $O(1)$ Landau-Zener correction, explained in Sec. 7.1.2, are visible. We believe that the adiabatic numerics ($v = 0$ red dot) differ from the analytics (blue line) due to finite-size effects: For small systems away from the spectrum's center, the average gap estimated from the density of states can vary appreciably over one gap. These numerics confirm the analytics and signal the MBL Otto engine's robustness with respect to changes in the tuning protocol.




Figure 7.3: Level-spacing distribution for 10^6 disorder realizations of the random-field Heisenberg model at disorder width $h = 20$ and system size $N = 8$ (blue line). The vertical black line shows the estimate of the level-repulsion parameter δ_- .

7.2.2 Enlightenment from numerical methods

Sec. 7.2 discusses numerical simulations of a finite-speed MBL engine. We simulate not the notionally smooth tuning of Sec. 6.2, but rather a stepwise tuning with step size $\delta t \simeq 0.405 \langle \delta \rangle^{-1}$.

Faithfully simulating the smooth-tuning protocol is substantially more difficult than simulating the stepwise protocol. In either case, the tuning time grows exponentially with system size, because it is set by (an estimate of) the smallest gap. But in the stepwise protocol, we fix the step size in terms of the average gap, so the step size grows with the the tuning time.

Faithfully simulating the smooth tuning protocol, on the other hand, requires step size constant with system size. The highest frequency ω_{\max} in the problem is of order the many-body bandwidth $\sim \frac{1}{Q}N \sim 1$, and even with higher-order methods we require step size $\delta t \sim \omega_{\max}^{-1}$: to see this heuristically, note that a fourth-order polynomial can approximate at most two periods of a sine.

And yet the stepwise protocol with long timestep $\delta t \simeq 0.405 \langle \delta \rangle^{-1}$ captures the essential physics. Why is this? What about that physics is amenable to capture by such a crude method²? A timestep $\delta t = 0.405 \langle \delta \rangle^{-1}$ can capture physics on energy scales smaller than $\delta t^{-1} \simeq 2 \langle \delta \rangle$. The results of Fig. 7.2 are

²Given our reliance on the numerics to catch persuasive but subtly incorrect arguments, this could be merely signal that we made a mistake that happened to line up with the ways in which the stepwise protocol differs from the smooth protocol. But let us proceed on the assumption that this is not the case.

not sensitive to δt , suggesting that diabatic corrections occur on energy scales small. This is (weak) evidence for our implicit supposition in Sec. 7.1.1: a diabatic correction that far-away levels (that is, levels that are not nearest in energy). It is also evidence that the fractional Landau-Zener and Landau-Zener transitions of Sec's 7.1.1 and 7.1.2—which take place on an energy scale δ_- —are indeed the dominant contribution.

To claim this definitively would require a much more careful analysis of convergence in δt than we have performed. Looking beyond convergence analysis, this is a place where the well-developed methods of numerical solution of ODEs may have physical insight to offer. In particular, adaptive stepsize methods would directly probe the energy scale required to capture the system's essential physics. The challenge in applying these methods would be constructing a good quantitative measure of how accurately an individual step has been approximated.

7.3 Precluding communication between subengines

To maintain the MBL engine's advantage, we must approximately isolate subengines. The subengines' (near) independence implies a lower bound on the tuning speed v : The price paid for scalability is the impossibility of adiabaticity. Suppose that $H_{\text{macro}}(t)$ were tuned infinitely slowly. Information would have time to propagate from one subengine to every other. The slow spread of information through MBL [109] lower-bounds v . This consideration, however, does not turn out to be the most restrictive constraint on the cycle time. Therefore, we address it only qualitatively.

As explained in Sec. 7.1.2, v determines the effective size of an MBL subengine. Ideally, v is large enough to prevent adiabatic transitions between configurations extended beyond the mesoscale N . For each stage of the engine's operation, v should exceed the speed given in Eq. (7.7) for the localization length ξ of a length- $(N + 1)$ chain:

$$v \gg [\delta_-(N + 1, \xi)]^2 \sim \mathcal{E}^2 2^{-2(N+1)} e^{-2(N+1)/\xi}. \quad (7.13)$$

(We have made explicit the dependence of the level-repulsion scale δ_- on the mesoscale-engine size N and on the localization length ξ .)

This scale quickly decays as the system is taken through stroke 1. This implies that the speed should interpolate between $[\delta_-(N + 1, \xi_>)]^2$ and $\frac{(W_b)^3}{\delta_-(N, \xi_<)}$ [from Ineq. (7.6)].

7.4 Lower bound on the cycle time from cold thermalization

Thermalization with the cold bath (stroke 2) bounds τ_{cycle} more stringently than the Hamiltonian tunings do. The reasons are (i) the slowness with which MBL thermalizes and (ii) the restriction $W_b \ll \langle \delta \rangle$ on the cold-bath bandwidth. We elaborate after introducing our cold-thermalization model (see [238, App. I] for details).

We envision the cold bath as a bosonic system that couples to the engine locally, as via the Hamiltonian

$$H_{\text{int}} = g \int_{-W_b/\xi}^{W_b/\xi} d\omega \sum_{j=1}^{N_{\text{macro}}} \left(c_j^\dagger c_{j+1} + \text{h.c.} \right) (b_\omega + b_\omega^\dagger) \times \delta(\langle 0 | c_j H_{\text{macro}}(\tau) c_{j+1}^\dagger | 0 \rangle - \omega). \quad (7.14)$$

The coupling strength is denoted by g . c_j and c_j^\dagger denote the annihilation and creation of a fermion at site j . $H_{\text{macro}}(t)$ denotes the Hamiltonian that would govern the engine at time t in the bath's absence. Cold thermalization lasts from $t = \tau$ to $t = \tau'$ (Fig. 6.1). b_ω and b_ω^\dagger represent the annihilation and creation of a frequency- ω boson in the bath. The Dirac delta function is denoted by $\delta(\cdot)$.

The bath couples locally, e.g., to pairs of nearest-neighbor spins. This locality prevents subengines from interacting with each other much through the bath. The bath can, e.g., flip spin j upward while flipping spin $j + 1$ downward. These flips likely change a subengine's energy by an amount E . The bath can effectively absorb only energy quanta of size $\leq W_b$ from any subengine. The cap is set by the bath's speed of sound [114], which follows from microscopic parameters in the bath's Hamiltonian [134]. The rest of the energy emitted during the spin flips, $|E - W_b|$, is distributed across the subengine as the intrinsic subengine Hamiltonian flips more spins.

Let τ_{th} denote the time required for stroke 2. We estimate τ_{th} from Fermi's Golden Rule,

$$\Gamma_{fi} = \frac{2\pi}{\hbar} |\langle f | V | i \rangle|^2 \mu_{\text{bath}}. \quad (7.15)$$

Cold thermalization transitions the engine from an energy level $|i\rangle$ to a level $|f\rangle$. The bath has a density of states $\mu_{\text{bath}} \sim 1/W_b$.

We estimate the matrix-element size $|\langle f | V | i \rangle|$ as follows. Cold thermaliza-

tion transfers energy $E_{if} \sim W_b$ from the subengine to the bath. W_b is very small. Hence the energy change rearranges particles across a large distance $L \gg \xi = \xi_<$, due to local level correlations. V nontrivially transforms just a few subengine sites. Such a local operator rearranges particles across a large distance L at a rate that scales as $\mathcal{E}e^{-L/\xi} 2^{-L} \sim \delta_-$. Whereas \mathcal{E} sets the scale of the level repulsion δ_- , g sets the scale of $|\langle f|V|i\rangle|$. The correlation length $\xi = \xi_<$ during cold thermalization. We approximate L with the subengine length $\xi_>$. Hence $|\langle f|V|i\rangle| \sim \frac{g\delta_-}{\mathcal{E}}$.

We substitute into Eq. (7.15). The transition rate $\Gamma_{fi} = \frac{1}{\tau_{\text{th}}}$. Inverting yields

$$\tau_{\text{cycle}} \sim \tau_{\text{th}} \sim W_b \left(\frac{\mathcal{E}}{g\delta_-} \right)^2. \quad (7.16)$$

To bound τ_{cycle} , we must bound the coupling g . The interaction is assumed to be Markovian: Information leaked from the engine dissipates throughout the bath quickly. Bath correlation functions must decay much more quickly than the coupling transfers energy. If τ_{bath} denotes the correlation-decay time, $\tau_{\text{bath}} < \frac{1}{g}$. The small-bandwidth bath's $\tau_{\text{bath}} \sim 1/W_b$, so $g < W_b$. This inequality, with Ineq. (7.16), implies

$$\tau_{\text{cycle}} = \tau_{\text{th}} > \frac{\mathcal{E}^2}{W_b(\delta_-)^2} \sim \frac{10}{\mathcal{E}} e^{2\xi_>/\xi_<} 2^{3\xi_>}. \quad (7.17)$$

The final expression follows if $W_b \sim \frac{\langle \delta \rangle}{10}$.

Like Markovianity, higher-order processes bound τ_{th} . Such processes transfer energy $E > W_b$ between the engine and the cold bath. Because they require $n \sim E/W_b$ excitations in the bath, they will be suppressed by a factor $\sim g^{E/W_b}$.

To understand these processes in more detail, write $H_{\text{tot}} = H_{\text{macro}}(t) + H_{\text{bath}} + H_{\text{int}}$ for the Hamiltonian that governs the engine-and-bath composite. H_{tot} generates the time-evolution operator $U(t) := e^{-iH_{\text{tot}}t}$. Consider Taylor-expanding $U(t)$. The ℓ th term is suppressed in g^ℓ , contains 2ℓ fermion operators c_j and c_j^\dagger , and contains ℓ boson operators b_ω and b_ω^\dagger . This term encodes the absorption, by the bath, of ℓ energy quanta of sizes $\leq W_b$. The subengine gives the bath a total amount $\sim \ell W_b$ of heat. The subengine should not lose so much heat. Hence higher-order processes should occur much more slowly than the

rate- g processes:

$$\tau_{\text{high-ord.}} \gg \tau_{\text{th}}. \quad (7.18)$$

Let us construct an expression for the left-hand side. Which processes most urgently require suppressing? Processes that change the subengine's energy by $\gtrsim \langle \delta \rangle$. Figure 6.2 illustrates why. If the right-hand leg has length $\gtrsim \langle \delta \rangle$, the right-hand leg could be longer than the left-hand leg. If it were, the trial would yield net negative work, $W_{\text{tot}} < 0$. The bath would absorb energy $\langle \delta \rangle$ from a subengine by absorbing $\sim \frac{\langle \delta \rangle}{W_b}$ packets of energy $\sim W_b$ each. Hence the bath would appear to need to flip $\sim L = \frac{\langle \delta \rangle}{W_b}$ spins to absorb energy $\sim \langle \delta \rangle$. (We switch from fermion language to spin language for convenience.) However, the length- L spin subchain has a discrete effective energy spectrum. The spectrum might lack a level associated with the amount (initial energy) $-\langle \delta \rangle$ of energy. If so, the bath must flip more than $\frac{\langle \delta \rangle}{W_b}$ spins—local level correlations suggest $\sim \xi_>$ spins. Hence $L = \max \left\{ \frac{\langle \delta \rangle}{W_b}, \xi_> \right\}$. Energy is rearranged across the distance L at a rate $\propto g^L$.

RACING THE MBL-MOBILE

In Chs 7, we analyzed the MBL engine mostly in terms of scaling properties. But what are the constant factors in front of those scalings? In particular, what would the power density of an experimental realization look like? And how does our engine compare to other quantum Otto engines?

8.1 Comparisons with competitor Otto engines

We compare the MBL Otto engine to four competitors: a bandwidth engine, a variant of the MBL engine that is tuned between two disorder strengths, an engine of quantum dots, and an Anderson-localized engine. We argue that the MBL Otto engine is more robust against perturbations than the bandwidth, Anderson, and quantum-dot engines. We also argue that our MBL engine is more reliable than the equal-disorder-strength engine: Our MBL engine’s W_{tot} varies less from trial to trial and suppresses worst-case trials, in which $W_{\text{tot}} < 0$. This paper’s arguments go through almost unchanged for an Anderson-localized medium. Such a medium would lack robustness against interactions, though: Even if the interactions do not delocalize the medium—which would destroy the engine—they would turn the Anderson engine into an MBL engine. One can view our MBL engine as an easy generalization of the Anderson engine.

8.1.1 Comparison with bandwidth engine

Imagine eliminating the scaling factor $Q(h(\alpha_t))$ from the Hamiltonian (6.33). The energy band is compressed and expanded as the disorder strength $h(\alpha_t)$ is ramped down and up. The whole band, rather than a gap, contracts and widens as in Fig. 6.2, between a size $\sim \mathcal{E}N_{\text{macro}} h(\alpha_0)$ and a size $\sim \mathcal{E}N_{\text{macro}} h(\alpha_1) \gg \mathcal{E}N_{\text{macro}} h(\alpha_0)$. The engine can remain in one phase throughout the cycle. The cycle does not benefit from the “athermality” of local level correlations.

Furthermore, this accordion-like motion requires no change of the energy eigenbasis’s form. Tuning may proceed quantum-adiabatically: $v \approx 0$. The ideal engine suffers no diabatic jumps, losing $\langle W_{\text{diab}} \rangle_{\text{macro}} = 0$.

But this engine is impractical: Consider any perturbation V that fails to commute with the ideal Hamiltonian $H(t)$: $[V, H(t)] \neq 0$. Stray fields, for example, can taint an environment. As another example, consider cold atoms in an optical lattice. The disorder strength is ideally $\mathcal{E}h(\alpha_t)$. One can strengthen the disorder by strengthening the lattice potential U_{lattice} . Similarly, one can raise the hopping frequency (ideally \mathcal{E}) by raising the pressure p . Strengthening U_{lattice} and p while achieving the ideal disorder-to-hopping ratio $\frac{\mathcal{E}h(\alpha_t)}{\mathcal{E}} = h(\alpha_t)$ requires fine control. If the ratio changes from $h(\alpha_t)$, the Hamiltonian $H(t)$ acquires a perturbation V that fails to commute with other terms.

This V can cause diabatic jumps that cost work $\langle W_{\text{diab}} \rangle_{\text{macro}}$. Can the bandwidth engine not withstand several hops—say, through $0.02\mathcal{N}_{\text{macro}}$ levels?

No, because the ground state pulls away from the rest of the spectrum as N_{macro} grows. Suppose, for simplicity, that $T_C = 0$ and $T_H = \infty$. The bandwidth engine starts stroke 1 in $\rho(0) = I/\mathcal{N}_{\text{macro}}$. Diabatic hops preserve $\rho(t)$ during stroke 1, on average: The engine as likely hops upward as drops. Cold thermalization drops the engine to the ground state (plus an exponentially small dusting of higher-level states). The ground-state energy is generically extensive. Hence the engine absorbs $\langle Q_2 \rangle_{\text{macro}} \sim -N_{\text{macro}}$, on average. Suppose that, during stroke 3, the engine jumps up through 2% of the levels. The engine ends about two standard deviations below the spectrum's center, with average energy $\sim \sqrt{N_{\text{macro}}}$. While returning to $T_H = 0$ during the average stroke 4, the bandwidth engine absorbs $\langle Q_4 \rangle_{\text{macro}} \sim \sqrt{N_{\text{macro}}}$. The average outputted work $\langle W_{\text{tot}} \rangle_{\text{macro}} = \langle Q_4 \rangle_{\text{macro}} + \langle Q_2 \rangle_{\text{macro}} \sim \sqrt{N_{\text{macro}}} - N_{\text{macro}}$. As N_{macro} grows, $\langle W_{\text{tot}} \rangle_{\text{macro}}$ dips farther below zero. A few diabatic jumps threaten the bandwidth engine's ability to output $\langle W_{\text{tot}} \rangle > 0$.

The bandwidth engine's v must decline as N_{macro} grows also because the typical whole-system gap $\langle \delta \rangle_{\text{macro}} \sim \frac{\mathcal{E}}{\mathcal{N}_{\text{macro}}}$ shrinks. The smaller the gaps, the greater the likelihood that a given v induces hops. As $\langle \delta \rangle_{\text{macro}} \rightarrow 0$, v must $\rightarrow 0$. The MBL Otto cycle proceeds more quickly, due to subengines' parallelization.

8.1.2 Comparison with MBL engine tuned between same-strength disorder realizations

Take our MBL Otto cycle, and vary not the disorder strength, but the disorder realization during each cycle. The disorder strength $h(\alpha_t)$ in Eq. (6.33) would remain $\gg 1$ and constant in t , while the random variables h_j would change. Let \tilde{S} denote this constant- $h(\alpha_t)$ engine, and let S denote the MBL engine.

\tilde{S} takes less advantage of MBL’s “athermality,” as \tilde{S} is not tuned between level-repelling and level-repulsion-free regimes.

Yet \tilde{S} outputs the amount $\langle W_{\text{tot}} \rangle$ of work outputted by S per cycle, on average. Because W_b is small, cold thermalization drops \tilde{S} across only small gaps $\delta' \ll \langle \delta \rangle$. \tilde{S} traverses a trapezoid, as in Fig. 6.2, in each trial. However, the MBL engine has two advantages: greater reliability and fewer worst-case (negative-work-outputted) trials.

Both the left-hand gap δ and the right-hand gap δ' traversed by \tilde{S} are Poisson-distributed. Poisson-distributed gaps more likely assume extreme values than GOE-distributed gaps: $P_{\text{MBL}}^{(E)}(\delta) > P_{\text{GOE}}^{(E)}(\delta)$ if $\delta \sim 0$ or $\delta \gg \langle \delta \rangle$ [44]. The left-hand gap δ traversed by S is GOE-distributed. Hence the W_{tot} outputted by \tilde{S} more likely assumes extreme values than the W_{tot} outputted by S . The greater reliability of S may suit S better to “one-shot statistical mechanics” [8, 25, 42, 43, 52, 55, 74, 75, 83, 91, 97, 210]. In one-shot theory, predictability of the work W_{tot} extractable in any given trial serves as a resource.

Additionally, S suffers fewer worst-case trials than \tilde{S} . We define as *worst-case* a trial in which the engine outputs net negative work, $W_{\text{tot}} < 0$. Consider again Fig. 6.2. Consider a similar figure that depicts the trapezoid traversed by \tilde{S} in some trial. The left-hand gap, δ , is distributed as the right-hand gap, δ' , is, according to $P_{\text{MBL}}^{(E)}(\delta)$. Hence δ has a decent chance of being smaller than δ' : $\delta < \delta'$. \tilde{S} would output $W_{\text{tot}} < 0$ in such a trial.

Suppose, for simplicity, that $T_H = \infty$ and $T_C = 0$. The probability that any given S trial outputs $W_{\text{tot}} < 0$ is

$$p_{\text{worst}} \approx (\text{Prob. that the left-hand gap} < \text{the right-hand gap}) \quad (8.1)$$

$$\times (\text{Prob. that the right-hand gap is small enough to be cold-thermalized})$$

$$\approx (\text{Prob. that the left-hand gap} < W_b) \times \frac{W_b}{\langle \delta \rangle}. \quad (8.2)$$

The initial factor is modeled by the area of a region under the $P_{\text{GOE}}^{(E)}(\delta)$ curve. The region stretches from $\delta = 0$ to $\delta = W_b$. We approximate the region as a triangle of length W_b and height $\frac{\pi}{2} \frac{W_b}{\langle \delta \rangle^2} e^{-\frac{\pi}{4} (W_b)^2 / \langle \delta \rangle^2} \sim \frac{W_b}{\langle \delta \rangle^2}$, [$\delta \approx W_b$, Eq. (5.2), and $\frac{W_b}{\langle \delta \rangle} \ll 1$]. The triangle has an area of $\frac{1}{2} \cdot W_b \cdot \frac{\pi}{2} \frac{W_b}{\langle \delta \rangle^2} \sim \left(\frac{W_b}{\langle \delta \rangle} \right)^2$. Substituting

into Eq. (8.2) yields

$$p_{\text{worst}} \sim \left(\frac{W_{\text{b}}}{\langle \delta \rangle} \right)^3. \quad (8.3)$$

Let \tilde{p}_{worst} denote the probability that any given \tilde{S} trial outputs $W_{\text{tot}} < 0$. \tilde{p}_{worst} shares the form of Eq. (8.2). The initial factor approximates to the area of a region under the $P_{\text{MBL}}^{(E)}(\delta)$ curve. The region extends from $\delta = 0$ to $\delta = W_{\text{b}}$. The region resembles a rectangle of height $P_{\text{MBL}}^{(E)}(0) \approx \frac{1}{\langle \delta \rangle}$. Combining the rectangle's area, $\frac{W_{\text{b}}}{\langle \delta \rangle}$, with Eq. (8.2) yields

$$\tilde{p}_{\text{worst}} \sim \left(\frac{W_{\text{b}}}{\langle \delta \rangle} \right)^2. \quad (8.4)$$

Since $\frac{W_{\text{b}}}{\langle \delta \rangle} \ll 1$, $p_{\text{worst}} \ll \tilde{p}_{\text{worst}}$.¹

8.1.3 Quantum-dot engine

MBL has been modeled with quasilocal bits [34, 93]. A string of ideally independent bits or qubits, such as quantum dots, forms a natural competitor. A qubit Otto engine's gap is shrunk, widened, and shrunk [11, 61, 67, 68, 88]. In addition to the order-of-magnitude analysis below, we make two points about implementations' practicality. First, the MBL potential's generic nature offers an advantage. MBL requires a random disorder potential $\{h(\alpha_t)h_j\}$, e.g., a “dirty sample,” a defect-riddled crystal. This “generic” potential contrasts with the pristine background required by quantum dots. Imposing random MBL disorder is expected to be simpler. On the other hand, a quantum-dot engine does not necessarily need a small-bandwidth cold bath, $W_{\text{b}} \ll \langle \delta \rangle$.

8.1.4 Anderson-localized engine

Anderson localization follows from removing the interactions from MBL. One could implement our Otto cycle with an Anderson insulator because Anderson Hamiltonians exhibit Poissonian level statistics (5.1). But strokes 1 and 3 would require the switching off and on of interactions. Tuning the interaction, as well as the disorder-to-interaction ratio, requires more effort than tuning just the latter.

¹The discrepancy is exaggerated if the exponent in Eq. (8.3) rises, if the left-hand S Hamiltonian is modeled with a Gaussian ensemble other than the GOE. The Gaussian unitary ensemble (GUE) contains an exponent of 4; the Gaussian symplectic ensemble (GSE), an exponent of 6. Different ensembles model different symmetries.

Also, particles typically interact in many-body systems. MBL particles interact; Anderson-localized particles do not. Hence one might eventually expect less difficulty in engineering MBL engines than in engineering Anderson-localized engines.

8.2 Order-of-magnitude estimates

Localized engine: Localization has been achieved in solid-state systems.² Consider silicon doped with phosphorus [124]. A distance of ~ 10 nm may separate phosphorus impurities. Let our engine cycle's shallowly localized regime have a localization length of $\xi_{>} \sim 10$ sites, or 100 nm. The work-outputting degrees of freedom will be electronic. The localized states will correspond to energies $\mathcal{E} \sim 1$ eV. Each subengine's half-filling Hilbert space has dimensionality $\mathcal{N} = \binom{10}{5} \sim 10^2$. Hence each subengine has an effective average gap $\langle \delta \rangle \sim \frac{\mathcal{E}\sqrt{\mathcal{N}}}{\mathcal{N}} \sim \frac{1\text{ eV}}{10^2} \sim 10$ meV. The cold-bath bandwidth must satisfy $\langle \delta \rangle \gg W_b$. We set W_b to be an order of magnitude down from $\langle \delta \rangle$: $W_b \sim 1$ meV ~ 10 K. The cold-bath bandwidth approximates the work outputted by one subengine per cycle:³ $\langle W_{\text{tot}} \rangle \sim W_b \sim 1$ meV [Eq. (6.7)].

What volume does a localized subengine fill? Suppose that the engine is three-dimensional (3D).⁴ A little room should separate the subengines. Classical-control equipment requires more room. Also, the subengine needs space to connect to the baths. We therefore associate each subengine with a volume of $V \approx (100 \text{ nm})^3$.

The last element needed is the cycle time, τ_{cycle} . We choose for δ_- to be a little smaller than W_b —of the same order: $\delta_- \sim W_b \sim 1$ meV. In the extreme case allowed by Ineq. (7.17), $\tau_{\text{cycle}} \sim \frac{\hbar\mathcal{E}^2}{W_b(\delta_-)^2} \sim \frac{\hbar\mathcal{E}^2}{(W_b)^3} \sim \frac{(10^{-15} \text{ eV s})(1 \text{ eV})^2}{(1 \text{ meV})^3} \sim 1 \mu\text{s}$.

The localized engine therefore operates with a power $\mathcal{P} \sim \frac{W_b}{\tau_{\text{cycle}}} \sim \frac{1 \text{ meV}}{1 \mu\text{s}} \approx 10^{-16}$ W. Interestingly, this \mathcal{P} is one order of magnitude greater than a flagellar motor's [27] power, according to our estimates.

²This localization is single-particle, or Anderson [12], rather than many-body. Sec. 8.1.4 extends the MBL Otto engine to an Anderson-localized Otto engine.

³The use of semiconductors would require corrections to our results. (Dipolar interactions would couple the impurities' spins. Energy eigenfunctions would decay as power laws with distance.) But we aim for just a rough estimate.

⁴Until now, we have supposed that the engine is 1D. Anderson localization, which has been realized in semiconductors, exists in all dimensionalities. Yet whether MBL exists in dimensionalities $D > 1$ remains an open question. Some evidence suggests that MBL exists in $D \geq 2$ [24, 37, 125]. But attributing a 3D volume to the engine facilitates comparisons with competitors. We imagine 10-nm-long 1D strings of sites. Strings are arrayed in a plane, separated by 10 nm. Planes are stacked atop each other, separated by another 10 nm.

We can assess the engine by calculating not only its power, but also its power density. The localized engine packs a punch at $\frac{\mathcal{P}}{V} \sim \frac{10^{-16} \text{ W}}{(10^{-7} \text{ m})^3} = 100 \text{ kW/m}^3$.

Car engine: The quintessential Otto engine powers cars. A typical car engine outputs $\mathcal{P} \sim 100$ horsepower $\sim 100 \text{ kW}$. A car’s power density is $\frac{\mathcal{P}}{V} \sim \frac{100 \text{ kW}}{100 \text{ L}} = 1 \text{ MW/m}^3$ (wherein L represents liters). The car engine’s $\frac{\mathcal{P}}{V}$ exceeds the MBL engine’s by only an order of magnitude, according to these rough estimates.

Array of quantum dots: MBL has been modeled with quasilocal bits [34, 93]. A string of ideally independent bits or qubits, such as quantum dots, forms a natural competitor. A qubit Otto engine’s gap is shrunk, widened, and shrunk [11, 61, 67, 68, 88].

A realization could consist of double quantum dots [165, 166]. The scales in [165, 166] suggest that a quantum-dot engine could output an amount $W_{\text{tot}} \sim 10 \text{ meV}$ of work per cycle per dot. We approximate the cycle time τ_{cycle} with the spin relaxation time: $\tau_{\text{cycle}} \sim 1 \mu\text{s}$. (The energy eigenbasis need not rotate, unlike for the MBL engine. Hence diabatic hops do not lower-bound the ideal-quantum-dot τ_{cycle} .) The power would be $\mathcal{P} \sim \frac{W_{\text{tot}}}{\tau_{\text{cycle}}} \sim \frac{10 \text{ meV}}{1 \mu\text{s}} \sim 10^{-15} \text{ W}$. The quantum-dot engine’s power exceeds the MBL engine’s by an order of magnitude.

However, the quantum dots must be separated widely. Otherwise, they will interact, as an ETH system. (See [122] for disadvantages of interactions in another quantum thermal machine. Spin-spin couplings cause “quantum friction,” limiting the temperatures to which a refrigerator can cool.) We compensate by attributing a volume $V \sim (1 \mu\text{m})^3$ to each dot. The power density becomes $\frac{\mathcal{P}}{V} \sim 1 \text{ kW/m}^3$, two orders of magnitude less than the localized engine’s. Localization naturally implies near independence of the subengines.

8.3 Outlook

The realization of thermodynamic cycles with quantum many-body systems was proposed very recently [32, 62, 99, 130, 149, 153, 164, 211]. MBL offers a natural platform, due to its “athermality” and to athermality’s resourcefulness in thermodynamics. We designed an Otto engine that benefits from the discrepancy between many-body-localized and “thermal” level statistics. The engine illustrates how MBL can be used for thermodynamic advantage.

Realizing the engine may provide a near-term challenge for existing experimental set-ups. Possible platforms include cold atoms [24, 37, 120, 148,

188]; nitrogen-vacancy centers [125]; ion traps [198]; and doped semiconductors [124], for which we provided order-of-magnitude estimates. Realizations will require platform-dependent corrections due to, e.g., variable-range hopping induced by particle-phonon interactions. As another example, semiconductors' impurities suffer from dipolar interactions. The interactions extend particles' wave functions from decaying exponentially across space to decaying as power laws.

Reversing the engine may pump heat from the cold bath to the hot, lowering the cold bath's temperature. Low temperatures facilitate quantum computation and low-temperature experiments. An MBL engine cycle might therefore facilitate state preparation and coherence preservation in quantum many-body experiments: A quantum many-body engine would cool quantum many-body systems.

We have defined as work the energy outputted during Hamiltonian tunings. Some battery must store this energy. We have refrained from specifying the battery's physical form, using an *implicit battery model*. An equivalent *explicit battery model* could depend on the experimental platform. Quantum-thermodynamics batteries have been modeled abstractly with ladder-like Hamiltonians [197]. An oscillator battery for our engine could manifest as the mode of an electromagnetic field in cavity quantum electrodynamics.

MBL is expected to have thermodynamic applications beyond this Otto engine. A localized ratchet, for example, could leverage information to transform heat into work. Additionally, the paucity of transport in MBL may have technological applications beyond thermodynamics. Dielectrics, for example, prevent particles from flowing in undesirable directions. But dielectrics break down in strong fields. To survive, a dielectric must insulate well—as does MBL.

In addition to suggesting applications of MBL, this work identifies an opportunity within quantum thermodynamics. Athermal quantum states (e.g., $\rho \neq e^{-H/T}/Z$) are usually regarded as resources in quantum thermodynamics [26, 43, 50, 69, 75, 77, 83, 91, 98, 140, 141, 229, 236, 237]. Not only athermal states, we have argued, but also athermal energy-level statistics, offer thermodynamic advantages. Generalizing the quantum-thermodynamics definition of “resource” may expand the set of goals that thermodynamic agents can achieve.

Optimization offers another theoretical opportunity. We have shown that the engine works, but better protocols could be designed. For example, we pre-

scribe nearly quantum-adiabatic tunings. Shortcuts to adiabaticity (STA) avoid both diabatic transitions and exponentially slow tunings [2, 35, 51, 53, 122, 209]. STA have been used to reduce other quantum engines' cycle times [2, 51, 53]. STA might be applied to the many-body Otto cycle, after being incorporated into MBL generally.

BIBLIOGRAPHY

- [1] *ITensor Library (version 2.0.11)* <http://itensor.org>.
- [2] O. Abah and E. Lutz. Performance of superadiabatic quantum machines. *ArXiv e-prints*, November 2016.
- [3] O. Abah, J. Roßnagel, G. Jacob, S. Deffner, F. Schmidt-Kaler, K. Singer, and E. Lutz. Single-ion heat engine at maximum power. *Phys. Rev. Lett.*, 109:203006, Nov 2012.
- [4] Dmitry Abanin, Wojciech De Roeck, Wen Wei Ho, and François Huveneers. A rigorous theory of many-body prethermalization for periodically driven and closed quantum systems. *Commun. Math. Phys.*, 354(3):809–827, Sep 2017.
- [5] Dmitry A. Abanin, Ehud Altman, Immanuel Bloch, and Maksym Serbyn. Ergodicity, Entanglement and Many-Body Localization. *arXiv:1804.11065 [cond-mat, physics:quant-ph]*, April 2018. arXiv: 1804.11065.
- [6] Dmitry A. Abanin, Wojciech De Roeck, Wen Wei Ho, and Francois Huveneers. Effective hamiltonians, prethermalization, and slow energy absorption in periodically driven many-body systems. *Phys. Rev. B*, 95:014112, Jan 2017.
- [7] Dmitry A. Abanin, Wojciech De Roeck, and François Huveneers. Exponentially slow heating in periodically driven many-body systems. *Phys. Rev. Lett.*, 115:256803, Dec 2015.
- [8] Johan Åberg. Truly work-like work extraction via a single-shot analysis. *Nat. Commun.*, 4:1925, 2013.
- [9] Ian Affleck, Tom Kennedy, Elliott H. Lieb, and Hal Tasaki. Rigorous results on valence-bond ground states in antiferromagnets. *Physical Review Letters*, 59(7):799–802, August 1987.
- [10] Ian Affleck, Tom Kennedy, Elliott H. Lieb, and Hal Tasaki. Valence bond ground states in isotropic quantum antiferromagnets. *Communications in Mathematical Physics*, 115(3):477–528, September 1988.
- [11] G. Alvarado Barrios, F. Albarrán-Arriagada, F. A. Cárdenas-López, G. Romero, and J. C. Retamal. Role of quantum correlations in light-matter quantum heat engines. *ArXiv e-prints*, July 2017.
- [12] P. W. Anderson. Absence of diffusion in certain random lattices. *Phys. Rev.*, 109:1492–1505, Mar 1958.

- [13] M. C. Bañuls, J. I. Cirac, and M. B. Hastings. Strong and weak thermalization of infinite nonintegrable quantum systems. *Phys. Rev. Lett.*, 106:050405, Feb 2011.
- [14] M. C. Bañuls, M. B. Hastings, F. Verstraete, and J. I. Cirac. Matrix product states for dynamical simulation of infinite chains. *Phys. Rev. Lett.*, 102:240603, Jun 2009.
- [15] Satish Balay, William D. Gropp, Lois Curfman McInnes, and Barry F. Smith. Efficient management of parallelism in object oriented numerical software libraries. In E. Arge, A. M. Bruaset, and H. P. Langtangen, editors, *Modern Software Tools in Scientific Computing*, pages 163–202. Birkhäuser Press, 1997.
- [16] Jens H. Bardarson, Frank Pollmann, and Joel E. Moore. Unbounded Growth of Entanglement in Models of Many-Body Localization. *Physical Review Letters*, 109(1):017202, July 2012.
- [17] Thomas Barthel. Precise evaluation of thermal response functions by optimized density matrix renormalization group schemes. *New Journal of Physics*, 15(7):073010, 2013.
- [18] Thomas Barthel, Ulrich Schollwöck, and Steven R. White. Spectral functions in one-dimensional quantum systems at finite temperature using the density matrix renormalization group. *Phys. Rev. B*, 79:245101, Jun 2009.
- [19] D.M. Basko, I.L. Aleiner, and B.L. Altshuler. Metal–insulator transition in a weakly interacting many-electron system with localized single-particle states. *Annals of Physics*, 321(5):1126–1205, May 2006.
- [20] M. C. Bañuls, J. I. Cirac, and M. B. Hastings. Strong and Weak Thermalization of Infinite Nonintegrable Quantum Systems. *Physical Review Letters*, 106(5):050405, February 2011.
- [21] Giuliano Benenti, Giulio Casati, Tomaž Prosen, Davide Rossini, and Marko Žnidarič. Charge and spin transport in strongly correlated one-dimensional quantum systems driven far from equilibrium. *Physical Review B*, 80(3):035110, July 2009.
- [22] Felix C Binder, Sai Vinjanampathy, Kavan Modi, and John Goold. Quantacell: powerful charging of quantum batteries. *New Journal of Physics*, 17(7):075015, 2015.
- [23] Moritz Binder and Thomas Barthel. Minimally entangled typical thermal states versus matrix product purifications for the simulation of equilibrium states and time evolution. *Phys. Rev. B*, 92:125119, Sep 2015.

- [24] Pranjali Bordia, Henrik Lüschen, Sebastian Scherg, Sarang Gopalakrishnan, Michael Knap, Ulrich Schneider, and Immanuel Bloch. Probing slow relaxation and many-body localization in two-dimensional quasiperiodic systems. *Phys. Rev. X*, 7:041047, Nov 2017.
- [25] Fernando Brandão, Mischa Horodecki, Woods, Nelly Ng, Jonathan Oppenheim, and Stephanie Wehner. The second laws of quantum thermodynamics. 112(11):3275–3279, 2015.
- [26] Fernando G. S. L. Brandão, Michał Horodecki, Jonathan Oppenheim, Joseph M. Renes, and Robert W. Spekkens. Resource theory of quantum states out of thermal equilibrium. *Physical Review Letters*, 111(25):250404, December 2013.
- [27] Mostyn T. Brown. Bacterial flagellar motor: Biophysical studies. In Gordon C. K. Roberts, editor, *Encyclopedia of Biophysics*, pages 155–155. Springer Berlin Heidelberg, Berlin, Heidelberg, 2013.
- [28] Nicolas Brunner, Marcus Huber, Noah Linden, Sandu Popescu, Ralph Silva, and Paul Skrzypczyk. Entanglement enhances cooling in microscopic quantum refrigerators. *Phys. Rev. E*, 89:032115, Mar 2014.
- [29] Marin Bukov, Luca D’Alessio, and Anatoli Polkovnikov. Universal high-frequency behavior of periodically driven systems: from dynamical stabilization to floquet engineering. *Adv. Phys.*, 64(2):139–226, Mar 2015.
- [30] Marin Bukov, Sarang Gopalakrishnan, Michael Knap, and Eugene Demler. Prethermal floquet steady states and instabilities in the periodically driven, weakly interacting bose-hubbard model. *Phys. Rev. Lett.*, 115:205301, Nov 2015.
- [31] Selçuk Çakmak, Ferdi Altintas, Azmi Gençten, and Özgür E. Müstecaplıoğlu. Irreversible work and internal friction in a quantum otto cycle of a single arbitrary spin. *The European Physical Journal D*, 71(3):75, Mar 2017.
- [32] Michele Campisi and Rosario Fazio. The power of a critical heat engine. *Nature Communications*, 7:11895 EP –, Jun 2016. Article.
- [33] Amos Chan, Andrea De Luca, and J. T. Chalker. Spectral statistics in spatially extended chaotic quantum many-body systems. *Physical Review Letters*, 121(6), August 2018. arXiv: 1803.03841.
- [34] Anushya Chandran, Isaac H. Kim, Guifre Vidal, and Dmitry A. Abanin. Constructing local integrals of motion in the many-body localized phase. *Physical Review B*, 91(8):085425, February 2015. arXiv: 1407.8480.

- [35] Xi Chen, A. Ruschhaupt, S. Schmidt, A. del Campo, D. Guéry-Odelin, and J. G. Muga. Fast optimal frictionless atom cooling in harmonic traps: Shortcut to adiabaticity. *Phys. Rev. Lett.*, 104:063002, Feb 2010.
- [36] Xie Chen, Zheng-Cheng Gu, Zheng-Xin Liu, and Xiao-Gang Wen. Symmetry protected topological orders and the group cohomology of their symmetry group. *Physical Review B*, 87(15), April 2013.
- [37] J.-y. Choi, S. Hild, J. Zeiher, P. Schauß, A. Rubio-Abadal, T. Yefsah, V. Khemani, D. A. Huse, I. Bloch, and C. Gross. Exploring the many-body localization transition in two dimensions. *Science*, 352:1547–1552, June 2016.
- [38] Soonwon Choi, Joonhee Choi, Renate Landig, Georg Kucsko, Hengyun Zhou, Junichi Isoya, Fedor Jelezko, Shinobu Onoda, Hitoshi Sumiya, Vedika Khemani, Curt von Keyserlingk, Norman Y. Yao, Eugene Demler, and Mikhail D. Lukin. Observation of discrete time-crystalline order in a disordered dipolar many-body system. *Nature*, 543:221–225, Mar 2017.
- [39] Philippe Corboz. Variational optimization with infinite projected entangled-pair states. *Phys. Rev. B*, 94:035133, Jul 2016.
- [40] Gavin E. Crooks. Nonequilibrium measurements of free energy differences for microscopically reversible markovian systems. *Journal of Statistical Physics*, 90(5):1481–1487, 1998.
- [41] Jian Cui, J. Ignacio Cirac, and Mari Carmen Bañuls. Variational Matrix Product Operators for the Steady State of Dissipative Quantum Systems. *Physical Review Letters*, 114(22):220601, June 2015.
- [42] Oscar C. O. Dahlsten. Non-equilibrium statistical mechanics inspired by modern information theory. *Entropy*, 15(12):5346, 2013.
- [43] Oscar C. O. Dahlsten, Renato Renner, Elisabeth Rieper, and Vlatko Vedral. Inadequacy of von Neumann entropy for characterizing extractable work. *New J. Phys.*, 13(5):053015–1–053015–10, 2011.
- [44] Luca D’Alessio, Yariv Kafri, Anatoli Polkovnikov, and Marcos Rigol. From Quantum Chaos and Eigenstate Thermalization to Statistical Mechanics and Thermodynamics. *Adv. Phys.*, 65(3):239–362, May 2016.
- [45] Luca D’Alessio and Marcos Rigol. Long-time behavior of isolated periodically driven interacting lattice systems. *Phys. Rev. X*, 4:041048, Dec 2014.
- [46] C. De Grandi and A. Polkovnikov. Adiabatic Perturbation Theory: From Landau-Zener Problem to Quenching Through a Quantum Critical Point. In A. K. K. Chandra, A. Das, and B. K. K. Chakrabarti, editors,

Lecture Notes in Physics, Berlin Springer Verlag, volume 802 of *Lecture Notes in Physics, Berlin Springer Verlag*, page 75, 2010.

- [47] Andrea De Luca and Alberto Rosso. Dynamic nuclear polarization and the paradox of quantum thermalization. *Phys. Rev. Lett.*, 115:080401, Aug 2015.
- [48] Wojciech De Roeck, Francois Huveneers, Markus Müller, and Mauro Schiulaz. Absence of many-body mobility edges. *Phys. Rev. B*, 93:014203, Jan 2016.
- [49] Wojciech De Roeck and François Huveneers. Stability and instability towards delocalization in MBL systems. *Physical Review B*, 95(15), April 2017. arXiv: 1608.01815.
- [50] Sebastian Deffner, Juan Pablo Paz, and Wojciech H. Zurek. Quantum work and the thermodynamic cost of quantum measurements. *Phys. Rev. E*, 94:010103, Jul 2016.
- [51] A. del Campo, J. Goold, and M. Paternostro. More bang for your buck: Towards super-adiabatic quantum engines. *Scientific Reports*, 4, 2014.
- [52] Lidia del Río, Johan Aberg, Renato Renner, Oscar Dahlsten, and Vlatko Vedral. The thermodynamic meaning of negative entropy. *Nature*, 474(7349):61–63, Jun 2011.
- [53] Jiawen Deng, Qing-hai Wang, Zhihao Liu, Peter Hänggi, and Jiangbin Gong. Boosting work characteristics and overall heat-engine performance via shortcuts to adiabaticity: Quantum and classical systems. *Phys. Rev. E*, 88:062122, Dec 2013.
- [54] J. M. Deutsch. Quantum statistical mechanics in a closed system. *Phys. Rev. A*, 43:2046–2049, Feb 1991.
- [55] D Egloff, O C O Dahlsten, R Renner, and V Vedral. A measure of majorization emerging from single-shot statistical mechanics. *New Journal of Physics*, 17(7):073001, 2015.
- [56] Dominic V. Else, Bela Bauer, and Chetan Nayak. Floquet time crystals. *Phys. Rev. Lett.*, 117:090402, Aug 2016.
- [57] Dominic V. Else, Bela Bauer, and Chetan Nayak. Prethermal phases of matter protected by time-translation symmetry. *Phys. Rev. X*, 7:011026, Mar 2017.
- [58] Tilman Enss and Jesko Sirker. Light cone renormalization and quantum quenches in one-dimensional Hubbard models. *New Journal of Physics*, 14(2):023008, 2012.

- [59] M. Fannes, B. Nachtergaele, and R. F. Werner. Exact Antiferromagnetic Ground States of Quantum Spin Chains. *Europhysics Letters (EPL)*, 10(7):633–637, December 1989.
- [60] Adrian E. Feiguin and Gregory A. Fiete. Spectral properties of a spin-incoherent Luttinger liquid. *Physical Review B*, 81(7):075108, February 2010.
- [61] Tova Feldmann, Eitan Geva, Ronnie Kosloff, and Peter Salamon. Heat engines in finite time governed by master equations. *American Journal of Physics*, 64(4):485–492, 1996.
- [62] Dario Ferraro, Michele Campisi, Gian Marcello Andolina, Vittorio Pellegrini, and Marco Polini. High-power collective charging of a solid-state quantum battery. *Phys. Rev. Lett.*, 120:117702, Mar 2018.
- [63] Mark H Fischer, Mykola Maksymenko, and Ehud Altman. Dynamics of a many-body-localized system coupled to a bath. *Phys. Rev. Lett.*, 116:160401, Apr 2016.
- [64] Juan José García-Ripoll. Time evolution of matrix product states. *New Journal of Physics*, 8(12):305, 2006.
- [65] David Gelbwaser-Klimovsky and Alán Aspuru-Guzik. Strongly coupled quantum heat machines. *The Journal of Physical Chemistry Letters*, 6(17):3477–3482, 2015. PMID: 26291720.
- [66] J. E. Geusic, E. O. Schulz-DuBios, and H. E. D. Scovil. Quantum equivalent of the carnot cycle. *Phys. Rev.*, 156:343–351, Apr 1967.
- [67] Eitan Geva and Ronnie Kosloff. On the classical limit of quantum thermodynamics in finite time. *The Journal of Chemical Physics*, 97(6):4398–4412, 1992.
- [68] Eitan Geva and Ronnie Kosloff. A quantum-mechanical heat engine operating in finite time. a model consisting of spin-1/2 systems as the working fluid. *The Journal of Chemical Physics*, 96(4):3054–3067, 1992.
- [69] John Goold, Marcus Huber, Arnau Riera, Lúcia del Río, and Paul Skrzypczyk. The role of quantum information in thermodynamics: A topical review. *Journal of Physics A: Mathematical and Theoretical*, 49(14):143001, 2016.
- [70] Sarang Gopalakrishnan, Michael Knap, and Eugene Demler. Regimes of heating and dynamical response in driven many-body localized systems. *Phys. Rev. B*, 94:094201, Sep 2016.
- [71] Sarang Gopalakrishnan and Rahul Nandkishore. Mean-field theory of nearly many-body localized metals. *Phys. Rev. B*, 90:224203, Dec 2014.

- [72] Anna Goremykina, Romain Vasseur, and Maksym Serbyn. Analytically solvable renormalization group for the many-body localization transition. *arXiv:1807.04285 [cond-mat]*, July 2018. arXiv: 1807.04285.
- [73] Daniel Gottesman. The Heisenberg Representation of Quantum Computers. *arXiv:quant-ph/9807006*, July 1998. arXiv: quant-ph/9807006.
- [74] Gilad Gour. Quantum resource theories in the single-shot regime. *Phys. Rev. A*, 95:062314, Jun 2017.
- [75] Gilad Gour, Markus P. Müller, Varun Narasimhachar, Robert W. Spekkens, and Nicole Yunger Halpern. The resource theory of informational nonequilibrium in thermodynamics. *Physics Reports*, 583:1 – 58, 2015. The resource theory of informational nonequilibrium in thermodynamics.
- [76] T. Grover. Certain General Constraints on the Many-Body Localization Transition. *ArXiv e-prints*, May 2014.
- [77] Y. Guryanova, S. Popescu, A. J. Short, R. Silva, and P. Skrzypczyk. Thermodynamics of quantum systems with multiple conserved quantities. *Nature Communications*, 7:12049, July 2016.
- [78] Jutho Haegeman, J. Ignacio Cirac, Tobias J. Osborne, Iztok Pižorn, Henri Verschelde, and Frank Verstraete. Time-Dependent Variational Principle for Quantum Lattices. *Physical Review Letters*, 107(7):070601, August 2011.
- [79] Jutho Haegeman, Christian Lubich, Ivan Oseledets, Bart Vandereycken, and Frank Verstraete. Unifying time evolution and optimization with matrix product states. *Physical Review B*, 94(16):165116, October 2016.
- [80] Jutho Haegeman, Michaël Mariën, Tobias J. Osborne, and Frank Verstraete. Geometry of matrix product states: Metric, parallel transport, and curvature. *Journal of Mathematical Physics*, 55(2):021902, February 2014.
- [81] Jutho Haegeman, Tobias J. Osborne, and Frank Verstraete. Post-matrix product state methods: To tangent space and beyond. *Physical Review B*, 88(7):075133, August 2013.
- [82] Karen A. Hallberg. Density-matrix algorithm for the calculation of dynamical properties of low-dimensional systems. *Physical Review B*, 52(14):R9827–R9830, October 1995.
- [83] Nicole Yunger Halpern. Beyond heat baths ii: framework for generalized thermodynamic resource theories. *Journal of Physics A: Mathematical and Theoretical*, 51(9):094001, 2018.

- [84] Francis H Harlow and N Metropolis. Weapons Simulation Leads to the Computer Era. *Los Alamos Science*, page 10, 1983.
- [85] M. B. Hastings. Solving gapped hamiltonians locally. *Phys. Rev. B*, 73:085115, Feb 2006.
- [86] M. B. Hastings. An Area Law for One Dimensional Quantum Systems. *Journal of Statistical Mechanics: Theory and Experiment*, 2007(08):P08024–P08024, August 2007. arXiv: 0705.2024.
- [87] Johannes Hauschild and Frank Pollmann. Efficient numerical simulations with Tensor Networks: Tensor Network Python (TeNPy). *arXiv e-prints*, page arXiv:1805.00055, Apr 2018.
- [88] Jizhou He, Jincan Chen, and Ben Hua. Quantum refrigeration cycles using spin- $\frac{1}{2}$ systems as the working substance. *Phys. Rev. E*, 65:036145, Mar 2002.
- [89] Vicente Hernandez, Jose E. Roman, and Vicente Vidal. SLEPc: A scalable and flexible toolkit for the solution of eigenvalue problems. *ACM Trans. Math. Software*, 31(3):351–362, 2005.
- [90] Andreas Holzner, Andreas Weichselbaum, Ian P. McCulloch, Ulrich Schollwöck, and Jan von Delft. Chebyshev matrix product state approach for spectral functions. *Physical Review B*, 83(19):195115, May 2011.
- [91] Michał Horodecki and Jonathan Oppenheim. Fundamental limitations for quantum and nanoscale thermodynamics. *Nat. Commun.*, 4:1–6, 2013.
- [92] David A. Huse, Rahul Nandkishore, and Vadim Oganesyan. Phenomenology of fully many-body-localized systems. *Physical Review B*, 90(17):174202, November 2014.
- [93] David A. Huse, Rahul Nandkishore, and Vadim Oganesyan. Phenomenology of fully many-body-localized systems. *Phys. Rev. B*, 90:174202, Nov 2014.
- [94] David A. Huse, Rahul Nandkishore, Francesca Pietracaprina, Valentina Ros, and Antonello Scardicchio. Localized systems coupled to small baths: From anderson to zeno. *Phys. Rev. B*, 92:014203, Jul 2015.
- [95] Marko Žnidarič, Tomaž Prosen, and Peter Prelovšek. Many-body localization in the Heisenberg XXZ magnet in a random field. *Phys. Rev. B*, 77:064426, Feb 2008.
- [96] Yoseph Imry and Shang-keng Ma. Random-field instability of the ordered state of continuous symmetry. *Phys. Rev. Lett.*, 35:1399–1401, Nov 1975.

- [97] Kosuke Ito and Masahito Hayashi. Optimal performance of generalized heat engines with finite-size baths of arbitrary multiple conserved quantities beyond independent-and-identical-distribution scaling. *Phys. Rev. E*, 97:012129, Jan 2018.
- [98] Dominik Janzing, Pawel Wocjan, Robert Zeier, Rubino Geiss, and Th Beth. Thermodynamic cost of reliability and low temperatures: tightening Landauer’s principle and the Second Law. *Int. J. Theor. Phys.*, 39(12):2717–2753, 2000.
- [99] J. Jaramillo, M. Beau, and A. del Campo. Quantum supremacy of many-particle thermal machines. *New Journal of Physics*, 18(7):075019, 2016.
- [100] Eric Jeckelmann. Dynamical density-matrix renormalization-group method. *Physical Review B*, 66(4):045114, July 2002.
- [101] Simon Jesenko and Marko Žnidarič. Finite-temperature magnetization transport of the one-dimensional anisotropic Heisenberg model. *Physical Review B*, 84(17):174438, November 2011.
- [102] B. Karimi and J. P. Pekola. Otto refrigerator based on a superconducting qubit: Classical and quantum performance. *Phys. Rev. B.*, 94(18):184503, November 2016.
- [103] C. Karrasch, J. H. Bardarson, and J. E. Moore. Finite-temperature dynamical density matrix renormalization group and the drude weight of spin-1/2 chains. *Phys. Rev. Lett.*, 108:227206, May 2012.
- [104] C Karrasch, J H Bardarson, and J E Moore. Reducing the numerical effort of finite-temperature density matrix renormalization group calculations. *New Journal of Physics*, 15(8):083031, 2013.
- [105] C. Karrasch, D. M. Kennes, and F. Heidrich-Meisner. Spin and thermal conductivity of quantum spin chains and ladders. *Physical Review B*, 91(11):115130, March 2015.
- [106] D.M. Kennes and C. Karrasch. Extending the range of real time density matrix renormalization group simulations. *Computer Physics Communications*, 200:37 – 43, 2016.
- [107] Alexander V. Khaetskii, Daniel Loss, and Leonid Glazman. Electron spin decoherence in quantum dots due to interaction with nuclei. *Phys. Rev. Lett.*, 88:186802, Apr 2002.
- [108] Ilia Khait, Snir Gazit, Norman Y. Yao, and Assa Auerbach. Spin transport of weakly disordered heisenberg chain at infinite temperature. *Phys. Rev. B*, 93:224205, Jun 2016.

- [109] V. Khemani, R. Nandkishore, and S. L. Sondhi. Nonlocal adiabatic response of a localized system to local manipulations. *Nature Physics*, 11:560–565, July 2015.
- [110] Vedika Khemani, Achilleas Lazarides, Roderich Moessner, and S. L. Sondhi. Phase structure of driven quantum systems. *Phys. Rev. Lett.*, 116:250401, Jun 2016.
- [111] Vedika Khemani, Ashvin Vishwanath, and D. A. Huse. Operator spreading and the emergence of dissipation in unitary dynamics with conservation laws. *arXiv:1710.09835 [cond-mat, physics:hep-th, physics:quant-ph]*, October 2017. arXiv: 1710.09835.
- [112] Tien D. Kieu. The second law, maxwell’s demon, and work derivable from quantum heat engines. *Phys. Rev. Lett.*, 93:140403, Sep 2004.
- [113] Hyungwon Kim, Mari Carmen Bañuls, J. Ignacio Cirac, Matthew B. Hastings, and David A. Huse. Slowest local operators in quantum spin chains. *Physical Review E*, 92(1), July 2015.
- [114] Hyungwon Kim and David A. Huse. Ballistic spreading of entanglement in a diffusive nonintegrable system. *Phys. Rev. Lett.*, 111:127205, Sep 2013.
- [115] Hyungwon Kim, Tatsuhiko N. Ikeda, and David A. Huse. Testing whether all eigenstates obey the eigenstate thermalization hypothesis. *Phys. Rev. E*, 90:052105, Nov 2014.
- [116] Takuya Kitagawa, Erez Berg, Mark Rudner, and Eugene Demler. Topological characterization of periodically driven quantum systems. *Phys. Rev. B*, 82:235114, Dec 2010.
- [117] M. Kliesch, D. Gross, and J. Eisert. Matrix-product operators and states: Np-hardness and undecidability. *Phys. Rev. Lett.*, 113:160503, Oct 2014.
- [118] B. Kloss and Y. Bar Lev. Spin transport in long-range interacting one-dimensional chain. *ArXiv e-prints*, April 2018.
- [119] Benedikt Kloss, Yevgeny Bar Lev, and David Reichman. Time-dependent variational principle in matrix-product state manifolds: Pitfalls and potential. *Phys. Rev. B*, 97:024307, Jan 2018.
- [120] S. S. Kondov, W. R. McGehee, W. Xu, and B. DeMarco. Disorder-induced localization in a strongly correlated atomic hubbard gas. *Phys. Rev. Lett.*, 114:083002, Feb 2015.
- [121] Ronnie Kosloff and Tova Feldmann. Discrete four-stroke quantum heat engine exploring the origin of friction. *Phys. Rev. E*, 65:055102, May 2002.

- [122] Ronnie Kosloff and Tova Feldmann. Optimal performance of reciprocating demagnetization quantum refrigerators. *Phys. Rev. E*, 82:011134, Jul 2010.
- [123] Ronnie Kosloff and Yair Rezek. The quantum harmonic otto cycle. *Entropy*, 19(4):136, 2017.
- [124] B. Kramer and A. MacKinnon. Localization: theory and experiment. *Reports on Progress in Physics*, 56:1469–1564, December 1993.
- [125] G. Kucsko, S. Choi, J. Choi, P. C. Maurer, H. Sumiya, S. Onoda, J. Isoya, F. Jelezko, E. Demler, N. Y. Yao, and M. D. Lukin. Critical thermalization of a disordered dipolar spin system in diamond. *ArXiv e-prints*, September 2016.
- [126] Tomotaka Kuwahara, Takashi Mori, and Keiji Saito. Floquet-magnus theory and generic transient dynamics in periodically driven many-body quantum systems. *Ann. Phys.*, 367:96 – 124, 2016.
- [127] Till D. Kühner and Steven R. White. Dynamical correlation functions using the density matrix renormalization group. *Physical Review B*, 60(1):335–343, July 1999.
- [128] Achilleas Lazarides, Arnab Das, and Roderich Moessner. Equilibrium states of generic quantum systems subject to periodic driving. *Phys. Rev. E*, 90:012110, Jul 2014.
- [129] Pierre Le Doussal and Jonathan Machta. Annealed versus quenched diffusion coefficient in random media. *Phys. Rev. B*, 40:9427–9430, Nov 1989.
- [130] J. Lekscha, H. Wilming, J. Eisert, and R. Gallego. Quantum thermodynamics with local control. *ArXiv e-prints*, November 2016.
- [131] Marius Lemm. Gaplessness is not generic for translation-invariant spin chains. *arXiv:1903.00108 [cond-mat, physics:math-ph, physics:quant-ph]*, February 2019. arXiv: 1903.00108.
- [132] Emanuele Levi, Markus Heyl, Igor Lesanovsky, and Juan P. Garrahan. Robustness of many-body localization in the presence of dissipation. *Phys. Rev. Lett.*, 116:237203, Jun 2016.
- [133] E. Leviatan, F. Pollmann, J. H. Bardarson, and E. Altman. Quantum thermalization dynamics with Matrix-Product States. *ArXiv e-prints*, February 2017.
- [134] E. Lieb and D. Robinson. The finite group velocity of quantum spin systems. *Commun. Math. Phys.*, 28:251–257, 1972.

- [135] C.-J. Lin and O. I. Motrunich. Quasiparticle explanation of "weak thermalization" regime under quench in a non-integrable quantum spin chain. *ArXiv e-prints*, October 2016.
- [136] Cheng-Ju Lin and Olexei I. Motrunich. Explicit construction of quasi-conserved local operator of translationally invariant nonintegrable quantum spin chain in prethermalization. *Physical Review B*, 96(21):214301, December 2017.
- [137] Shi-Zeng Lin and Satoru Hayami. Ginzburg-landau theory for skyrmions in inversion-symmetric magnets with competing interactions. *Phys. Rev. B*, 93:064430, Feb 2016.
- [138] Netanel H. Lindner, Gil Refael, and Victor Galitski. Floquet Topological Insulator in Semiconductor Quantum Wells. *Nat. Phys.*, 7(6):490–495, June 2011.
- [139] Y.-K. Liu. Consistency of Local Density Matrices is QMA-complete. *eprint arXiv:quant-ph/0604166*, April 2006.
- [140] M. Lostaglio, D. Jennings, and T. Rudolph. Description of quantum coherence in thermodynamic processes requires constraints beyond free energy. *Nature Communications*, 6:6383, March 2015.
- [141] Matteo Lostaglio, David Jennings, and Terry Rudolph. Thermodynamic resource theories, non-commutativity and maximum entropy principles. *New Journal of Physics*, 19(4):043008, 2017.
- [142] Elihu Lubkin and Thelma Lubkin. Average quantal behavior and thermodynamic isolation. *Int. J. Theor. Phys.*, 32:933, 1993. Article.
- [143] D. J. Luitz and Y. Bar Lev. The Ergodic Side of the Many-Body Localization Transition. *ArXiv e-prints*, October 2016.
- [144] David J. Luitz and Yevgeny Bar Lev. Anomalous thermalization in ergodic systems. *Phys. Rev. Lett.*, 117:170404, Oct 2016.
- [145] David J. Luitz, François Huveneers, and Wojciech de Roeck. How a small quantum bath can thermalize long localized chains. *Physical Review Letters*, 119(15), October 2017. arXiv: 1705.10807.
- [146] David J. Luitz, Nicolas Laflorencie, and Fabien Alet. Many-body localization edge in the random-field heisenberg chain. *Phys. Rev. B*, 91:081103, Feb 2015.
- [147] David J. Luitz, Nicolas Laflorencie, and Fabien Alet. Extended slow dynamical regime close to the many-body localization transition. *Phys. Rev. B*, 93:060201, Feb 2016.

- [148] Henrik P. Lüschen, Pranjal Bordia, Sean S. Hodgman, Michael Schreiber, Saubhik Sarkar, Andrew J. Daley, Mark H. Fischer, Ehud Altman, Immanuel Bloch, and Ulrich Schneider. Signatures of many-body localization in a controlled open quantum system. *Phys. Rev. X*, 7:011034, Mar 2017.
- [149] Y.-H. Ma, S.-H. Su, and C.-P. Sun. Quantum thermodynamic cycle with quantum phase transition. *Phys. Rev. E*, 96(2):022143, August 2017.
- [150] Francisco Machado, Gregory D. Meyer, Dominic V. Else, Chetan Nayak, and Norman Y. Yao. Exponentially Slow Heating in Short and Long-range Interacting Floquet Systems. *arXiv:1708.01620*, August.
- [151] Ugo Marzolino and Tomaž Prosen. Computational complexity of nonequilibrium steady states of quantum spin chains. *Physical Review A*, 93(3):032306, March 2016.
- [152] Eduardo Mascarenhas, Hugo Flayac, and Vincenzo Savona. Matrix-product-operator approach to the nonequilibrium steady state of driven-dissipative quantum arrays. *Phys. Rev. A*, 92:022116, Aug 2015.
- [153] Ranjan Modak and Marcos Rigol. Work extraction in an isolated quantum lattice system: Grand canonical and generalized gibbs ensemble predictions. *Phys. Rev. E*, 95:062145, Jun 2017.
- [154] Andras Molnar, Norbert Schuch, Frank Verstraete, and J. Ignacio Cirac. Approximating gibbs states of local hamiltonians efficiently with projected entangled pair states. *Phys. Rev. B*, 91:045138, Jan 2015.
- [155] Takashi Mori, Tomotaka Kuwahara, and Keiji Saito. Rigorous bound on energy absorption and generic relaxation in periodically driven quantum systems. *Phys. Rev. Lett.*, 116:120401, Mar 2016.
- [156] Subroto Mukerjee, Vadim Oganesyan, and David Huse. Statistical theory of transport by strongly interacting lattice fermions. *Phys. Rev. B*, 73:035113, Jan 2006.
- [157] Alexander Müller-Hermes, J Ignacio Cirac, and Mari Carmen Bañuls. Tensor network techniques for the computation of dynamical observables in one-dimensional quantum spin systems. *New Journal of Physics*, 14(7):075003, 2012.
- [158] R. Nandkishore and D. A. Huse. Many-Body Localization and Thermalization in Quantum Statistical Mechanics. *Annual Review of Condensed Matter Physics*, 6:15–38, March 2015.
- [159] Vadim Oganesyan and David A. Huse. Localization of interacting fermions at high temperature. *Phys. Rev. B*, 75:155111, Apr 2007.

- [160] Arijeet Pal and David A. Huse. Many-body localization phase transition. *Physical Review B*, 82(17), November 2010.
- [161] S. A. Parameswaran and S. Gopalakrishnan. Spin-catalyzed hopping conductivity in disordered strongly interacting quantum wires. *Phys. Rev. B*, 95:024201, Jan 2017.
- [162] Daniel E. Parker, Xiangyu Cao, Thomas Scaffidi, and Ehud Altman. A universal operator growth hypothesis. *arXiv:1812.08657*, Dec.
- [163] D. Pekker and B. K. Clark. Encoding the structure of many-body localization with matrix product operators. *ArXiv e-prints*, October 2014.
- [164] Martí Perarnau-Llobet, Arnau Riera, Rodrigo Gallego, Henrik Wilming, and Jens Eisert. Work and entropy production in generalised gibbs ensembles. *New Journal of Physics*, 18(12):123035, 2016.
- [165] J. R. Petta, A. C. Johnson, J. M. Taylor, E. A. Laird, A. Yacoby, M. D. Lukin, C. M. Marcus, M. P. Hanson, and A. C. Gossard. Coherent manipulation of coupled electron spins in semiconductor quantum dots. *Science*, 309(5744):2180–2184, 2005.
- [166] J.R. Petta, A.C. Johnson, J.M. Taylor, A. Yacoby, M.D. Lukin, C.M. Marcus, M.P. Hanson, and A.C. Gossard. Charge and spin manipulation in a few-electron double dot. *Physica E: Low-dimensional Systems and Nanostructures*, 34(1-2):42 – 46, 2006. Proceedings of the 16th International Conference on Electronic Properties of Two-Dimensional Systems (EP2DS-16).
- [167] Iztok Pižorn, Viktor Eisler, Sabine Andergassen, and Matthias Troyer. Real time evolution at finite temperatures with operator space matrix product states. *New Journal of Physics*, 16(7):073007, 2014.
- [168] Iztok Pižorn and Tomaž Prosen. Operator space entanglement entropy in XY spin chains. *Physical Review B*, 79(18):184416, May 2009.
- [169] Frank Pollmann, Erez Berg, Ari M. Turner, and Masaki Oshikawa. Symmetry protection of topological phases in one-dimensional quantum spin systems. *Physical Review B*, 85(7):075125, February 2012.
- [170] Frank Pollmann, Ari M. Turner, Erez Berg, and Masaki Oshikawa. Entanglement spectrum of a topological phase in one dimension. *Physical Review B*, 81(6):064439, February 2010.
- [171] Ionut-Dragos Potirniche, Sumilan Banerjee, and Ehud Altman. On the stability of many-body localization in $d > 1$. *arXiv:1805.01475 [cond-mat]*, May 2018. arXiv: 1805.01475.

- [172] Andrew C. Potter, Romain Vasseur, and S. A. Parameswaran. Universal properties of many-body delocalization transitions. *Phys. Rev. X*, 5:031033, Sep 2015.
- [173] T. Prosen. Time evolution of a quantum many-body system: Transition from integrability to ergodicity in the thermodynamic limit. *Phys. Rev. Lett.*, 80:1808–1811, Mar 1998.
- [174] T. Prosen. Ergodic properties of a generic nonintegrable quantum many-body system in the thermodynamic limit. *Phys. Rev. E*, 60:3949–3968, Oct 1999.
- [175] Tomaž Prosen and Iztok Pižorn. Operator space entanglement entropy in a transverse Ising chain. *Physical Review A*, 76(3):032316, September 2007.
- [176] Tomaž Prosen and Iztok Pižorn. Quantum Phase Transition in a Far-from-Equilibrium Steady State of an XY Spin Chain. *Physical Review Letters*, 101(10):105701, September 2008.
- [177] Tomaž Prosen and Marko Žnidarič. Diffusive high-temperature transport in the one-dimensional Hubbard model. *Physical Review B*, 86(12):125118, September 2012.
- [178] Tomaž Prosen and Marko Žnidarič. Matrix product simulations of non-equilibrium steady states of quantum spin chains. *Journal of Statistical Mechanics: Theory and Experiment*, 2009(02):P02035, 2009.
- [179] Douglas Quattrochi. The internal combustion engine (otto cycle), Aug. 2006.
- [180] S. Ramasesha, Swapan K. Pati, H.R. Krishnamurthy, Z. Shuai, and J.L. Brédas. Low-lying electronic excitations and nonlinear optic properties of polymers via symmetrized density matrix renormalization group method. *Synthetic Metals*, 85(1-3):1019–1022, March 1997.
- [181] Peter Reimann. Typicality for generalized microcanonical ensembles. *Phys. Rev. Lett.*, 99:160404, Oct 2007.
- [182] Marcos Rigol, Vanja Dunjko, and Maxim Olshanii. Thermalization and its mechanism for generic isolated quantum systems. *Nature*, 452:854–858, April 2008.
- [183] Marcos Rigol, Vanja Dunjko, Vladimir Yurovsky, and Maxim Olshanii. Relaxation in a completely integrable many-body quantum system: An ab initio study of the dynamics of the highly excited states of 1d lattice hard-core bosons. *Phys. Rev. Lett.*, 98:050405, Feb 2007.

- [184] Brenden Roberts, Thomas Vidick, and Olexei I. Motrunich. Implementation of rigorous renormalization group method for ground space and low-energy states of local Hamiltonians. *Physical Review B*, 96(21):214203, December 2017. arXiv: 1703.01994.
- [185] J. E. Roman, C. Campos, E. Romero, and A. Tomas. SLEPc users manual. Technical Report DSIC-II/24/02 - Revision 3.7, D. Sistemes Informàtics i Computació, Universitat Politècnica de València, 2016.
- [186] Valerio Scarani, Mário Ziman, Peter Štelmachovič, Nicolas Gisin, and Vladimír Bužek. Thermalizing quantum machines: Dissipation and entanglement. *Phys. Rev. Lett.*, 88:097905, Feb 2002.
- [187] Ulrich Schollwöck. The density-matrix renormalization group in the age of matrix product states. *Annals of Physics*, 326(1):96 – 192, 2011. January 2011 Special Issue.
- [188] Michael Schreiber, Sean S. Hodgman, Pranjal Bordia, Henrik P. Lüschen, Mark H. Fischer, Ronen Vosk, Ehud Altman, Ulrich Schneider, and Immanuel Bloch. Observation of many-body localization of interacting fermions in a quasirandom optical lattice. *Science*, 349(6250):842–845, 2015.
- [189] Maximilian Schulz, Scott Richard Taylor, Christopher Andrew Hooley, and Antonello Scardicchio. Energy transport in a disordered spin chain with broken U(1) symmetry: diffusion, subdiffusion, and many-body localization. *arXiv:1805.01036 [cond-mat]*, May 2018. arXiv: 1805.01036.
- [190] Marlan O. Scully. Quantum afterburner: Improving the efficiency of an ideal heat engine. *Phys. Rev. Lett.*, 88:050602, Jan 2002.
- [191] Maksym Serbyn and Joel E. Moore. Spectral statistics across the many-body localization transition. *Phys. Rev. B*, 93:041424, Jan 2016.
- [192] Maksym Serbyn, Z. Papić, and Dmitry A. Abanin. Universal slow growth of entanglement in interacting strongly disordered systems. *Phys. Rev. Lett.*, 110:260601, Jun 2013.
- [193] Maksym Serbyn, Z. Papić, and Dmitry A. Abanin. Local Conservation Laws and the Structure of the Many-Body Localized States. *Physical Review Letters*, 111(12):127201, September 2013.
- [194] Maksym Serbyn, Z. Papić, and Dmitry A. Abanin. Universal Slow Growth of Entanglement in Interacting Strongly Disordered Systems. *Physical Review Letters*, 110(26):260601, June 2013.
- [195] S.N. Shevchenko, S. Ashhab, and Franco Nori. Landau-Zener-Stückelberg interferometry. *Physics Reports*, 492(1):1 – 30, 2010.

- [196] U. Sivan and Y. Imry. Energy-level correlation function and ac conductivity of a finite disordered system. *Phys. Rev. B*, 35:6074–6083, Apr 1987.
- [197] P. Skrzypczyk, A. J. Short, and S. Popescu. Extracting work from quantum systems. *ArXiv e-prints*, February 2013.
- [198] J. Smith, A. Lee, P. Richerme, B. Neyenhuis, P. W. Hess, P. Hauke, M. Heyl, D. A. Huse, and C. Monroe. Many-body localization in a quantum simulator with programmable random disorder. *Nat Phys*, 12(10):907–911, Oct 2016. Letter.
- [199] Qiao Song, Swati Singh, Keye Zhang, Weiping Zhang, and Pierre Meystre. One qubit and one photon: The simplest polaritonic heat engine. *Phys. Rev. A*, 94:063852, Dec 2016.
- [200] Mark Srednicki. Chaos and quantum thermalization. *Phys. Rev. E*, 50:888–901, Aug 1994.
- [201] S. V. Syzranov, A. V. Gorshkov, and V. Galitski. Out-of-time-order correlators in finite open systems. *ArXiv e-prints*, April 2017.
- [202] R. J. Sánchez, V. K. Varma, and V. Oganesyan. Anomalous and regular transport in spin- $\frac{1}{2}$ chains: ac conductivity. *Physical Review B*, 98(5):054415, August 2018.
- [203] Hiroshi Takasaki, Toshiya Hikihara, and Tomotoshi Nishino. Fixed Point of the Finite System DMRG. *Journal of the Physical Society of Japan*, 68(5):1537–1540, May 1999.
- [204] D. Tamascelli, R. Rosenbach, and M. B. Plenio. Improved scaling of time-evolving block-decimation algorithm through reduced-rank randomized singular value decomposition. *Physical Review E*, 91(6):063306, June 2015.
- [205] Hugo Terças, Sofia Ribeiro, Marco Pezzutto, and Yasser Omar. Quantum thermal machines driven by vacuum forces. *Phys. Rev. E*, 95:022135, Feb 2017.
- [206] Thimothée Thiery, François Huveneers, Markus Müller, and Wojciech De Roeck. Many-body delocalization as a quantum avalanche. *Physical Review Letters*, 121(14), October 2018. arXiv: 1706.09338.
- [207] Thimothée Thiery, Markus Müller, and Wojciech De Roeck. A microscopically motivated renormalization scheme for the MBL/ETH transition. *arXiv:1711.09880 [cond-mat]*, November 2017. arXiv: 1711.09880.
- [208] Alexander C. Tiegel, Salvatore R. Manmana, Thomas Pruschke, and Andreas Honecker. Matrix product state formulation of frequency-space dynamics at finite temperatures. *Phys. Rev. B*, 90:060406, Aug 2014.

- [209] E. Torrontegui, S. Ibáñez, S. Martínez-Garaot, M. Modugno, A. del Campo, D. Guéry-Odelin, A. Ruschhaupt, X. Chen, and J. G. Muga. Shortcuts to Adiabaticity. *Advances in Atomic Molecular and Optical Physics*, 62:117, 2013.
- [210] Remco van der Meer, Nelly Huei Ying Ng, and Stephanie Wehner. Smoothed generalized free energies for thermodynamics. *Phys. Rev. A*, 96:062135, Dec 2017.
- [211] Wouter Verstraelen, Dries Sels, and Michiel Wouters. Unitary work extraction from a generalized gibbs ensemble using bragg scattering. *Phys. Rev. A*, 96:023605, Aug 2017.
- [212] F. Verstraete, J. J. García-Ripoll, and J. I. Cirac. Matrix product density operators: Simulation of finite-temperature and dissipative systems. *Phys. Rev. Lett.*, 93:207204, Nov 2004.
- [213] Guifré Vidal. Efficient classical simulation of slightly entangled quantum computations. *Phys. Rev. Lett.*, 91:147902, Oct 2003.
- [214] Guifré Vidal. Efficient simulation of one-dimensional quantum many-body systems. *Phys. Rev. Lett.*, 93:040502, Jul 2004.
- [215] Sai Vinjanampathy and Janet Anders. Quantum thermodynamics. *Contemporary Physics*, 0(0):1–35, 0.
- [216] C. W. von Keyserlingk, Vedika Khemani, and S. L. Sondhi. Absolute stability and spatiotemporal long-range order in floquet systems. *Phys. Rev. B*, 94:085112, Aug 2016.
- [217] C. W. von Keyserlingk and S. L. Sondhi. Phase structure of one-dimensional interacting floquet systems. i. abelian symmetry-protected topological phases. *Phys. Rev. B*, 93:245145, Jun 2016.
- [218] C. W. von Keyserlingk and S. L. Sondhi. Phase structure of one-dimensional interacting floquet systems. ii. symmetry-broken phases. *Phys. Rev. B*, 93:245146, Jun 2016.
- [219] Curt von Keyserlingk, Tibor Rakovszky, Frank Pollmann, and Shivaji Sondhi. Operator hydrodynamics, OTOCs, and entanglement growth in systems without conservation laws. *Physical Review X*, 8(2), April 2018. arXiv: 1705.08910.
- [220] Ronen Vosk, David A. Huse, and Ehud Altman. Theory of the many-body localization transition in one-dimensional systems. *Phys. Rev. X*, 5:031032, Sep 2015.
- [221] Marko Žnidarič. Solvable quantum nonequilibrium model exhibiting a phase transition and a matrix product representation. *Physical Review E*, 83(1):011108, January 2011.

- [222] Marko Žnidarič, Tomaž Prosen, and Iztok Pižorn. Complexity of thermal states in quantum spin chains. *Physical Review A*, 78(2):022103, August 2008.
- [223] M L Wall and Lincoln D Carr. Out-of-equilibrium dynamics with matrix product states. *New Journal of Physics*, 14(12):125015, 2012.
- [224] Simon A. Weidinger and Michael Knap. Floquet prethermalization and regimes of heating in a periodically driven, interacting quantum system. *Sci. Rep.*, 7:45382, April 2017.
- [225] A. H. Werner, D. Jaschke, P. Silvi, M. Kliesch, T. Calarco, J. Eisert, and S. Montangero. Positive tensor network approach for simulating open quantum many-body systems. *Phys. Rev. Lett.*, 116:237201, Jun 2016.
- [226] Steven R. White. Density matrix formulation for quantum renormalization groups. *Physical Review Letters*, 69(19):2863–2866, November 1992.
- [227] Steven R. White. Density-matrix algorithms for quantum renormalization groups. *Physical Review B*, 48(14):10345–10356, October 1993.
- [228] Steven R. White and Adrian E. Feiguin. Real-time evolution using the density matrix renormalization group. *Phys. Rev. Lett.*, 93:076401, Aug 2004.
- [229] H. Wilming and R. Gallego. The third law as a single inequality. *ArXiv e-prints*, January 2017.
- [230] Kenneth G. Wilson. The renormalization group: Critical phenomena and the Kondo problem. *Reviews of Modern Physics*, 47(4):773–840, October 1975.
- [231] F. Alexander Wolf, Jorge A. Justiniano, Ian P. McCulloch, and Ulrich Schollwöck. Spectral functions and time evolution from the chebyshev recursion. *Phys. Rev. B*, 91:115144, Mar 2015.
- [232] M. P. Woods, N. Ng, and S. Wehner. The maximum efficiency of nano heat engines depends on more than temperature. *ArXiv e-prints*, June 2015.
- [233] Jonathan Wurtz and Anatoli Polkovnikov. Quantum Hydrodynamics in Spin Chains with Phase Space Methods. *ArXiv e-prints*, 1808:arXiv:1808.08977, August 2018.
- [234] Jonathan Wurtz, Anatoli Polkovnikov, and Dries Sels. Cluster Truncated Wigner Approximation in Strongly Interacting Systems. *Annals of Physics*, 395:341–365, August 2018. arXiv: 1804.10217.

- [235] N. Y. Yao, A. C. Potter, I.-D. Potirniche, and A. Vishwanath. Discrete time crystals: Rigidity, criticality, and realizations. *Phys. Rev. Lett.*, 118:030401, Jan 2017.
- [236] N. Yunger Halpern. Beyond heat baths II: Framework for generalized thermodynamic resource theories. *ArXiv e-prints*, September 2014.
- [237] N. Yunger Halpern, P. Faist, J. Oppenheim, and A. Winter. Microcanonical and resource-theoretic derivations of the thermal state of a quantum system with noncommuting charges. *Nature Communications*, 7:12051, July 2016.
- [238] N. Yunger Halpern, C. D. White, S. Gopalakrishnan, and G. Refael. MBL-mobile: Many-body-localized engine. *ArXiv e-prints*, July 2017.
- [239] Nicole Yunger Halpern, Andrew J. P. Garner, Oscar C. O. Dahlsten, and Vlatko Vedral. Introducing one-shot work into fluctuation relations. *New Journal of Physics*, 17(9):095003, 2015.
- [240] Michael P. Zaletel, Roger S. K. Mong, Christoph Karrasch, Joel E. Moore, and Frank Pollmann. Time-evolving a matrix product state with long-ranged interactions. *Physical Review B*, 91(16), April 2015. arXiv: 1407.1832.
- [241] Michael P. Zaletel, Roger S. K. Mong, Frank Pollmann, and Edward H. Rezayi. Infinite density matrix renormalization group for multicomponent quantum Hall systems. *Physical Review B*, 91(4), January 2015. arXiv: 1410.3861.
- [242] Bei Zeng, Xie Chen, Duan-Lu Zhou, and Xiao-Gang Wen. Quantum Information Meets Quantum Matter – From Quantum Entanglement to Topological Phase in Many-Body Systems. *arXiv:1508.02595 [cond-mat, physics:quant-ph]*, August 2015. arXiv: 1508.02595.
- [243] J. Zhang, P. W. Hess, A. Kyprianidis, P. Becker, A. Lee, J. Smith, G. Pagano, I.-D. Potirniche, A. C. Potter, A. Vishwanath, N. Y. Yao, and C. Monroe. Observation of a discrete time crystal. *Nature*, 543:217–220, Mar 2017.
- [244] Yuanjian Zheng and Dario Poletti. Work and efficiency of quantum otto cycles in power-law trapping potentials. *Phys. Rev. E*, 90:012145, Jul 2014.
- [245] M. Ziman, P. Stelmachovic, V. Buzek, M. Hillery, V. Scarani, and N. Gisin. Quantum homogenization. *eprint arXiv:quant-ph/0110164*, October 2001.

- [246] Marko Znidaric, Antonello Scardicchio, and Vipin Kerala Varma. Diffusive and subdiffusive spin transport in the ergodic phase of a many-body localizable system. *Physical Review Letters*, 117(4), July 2016. arXiv: 1604.08567.
- [247] Michael Zwolak and Guifré Vidal. Mixed-state dynamics in one-dimensional quantum lattice systems: A time-dependent superoperator renormalization algorithm. *Phys. Rev. Lett.*, 93:207205, Nov 2004.
- [248] Stellan Östlund and Stefan Rommer. Thermodynamic Limit of Density Matrix Renormalization. *Physical Review Letters*, 75(19):3537–3540, November 1995.

INDEX

- AKLT Hamiltonian, x
- Anderson localization, 2
- avalanche picture for MBL, 4
- bond dimension, 13
- canonical
 - mixed, x
- canonical form
 - left, 17
 - mixed, 19
 - right, 18
- convergence
 - false convergence, 93
- convergence testing, 140
- density matrix renormalization
 - group, *see* DMRG
- density matrix truncation, *see* DMT
- DMRG, ix
- DMT, xii
- eigenstate thermalization
 - hypothesis, xi
- ETH, *see* eigenstate thermalization hypothesis
- extended Church-Turing thesis, xi
- Floquet
 - drive, 59, 61
 - heating, 59
 - timescale, 61–64
- free fermion model, 93
- Gottesman-Knill theorem, xi
- isometry, 12
- l-bit, 2, 4
- many-body localization , *see* MBLxi
- matrix product operator , *see* MPO 21
- matrix product state, *see* MPSix,
see MPS11
- MBL, xi
- MPO, 21
- MPS, ix, 11
- orthogonality center, 19
- prethermalization
 - timescale, 61, 64
- prethermalization
 - Hamiltonian, 59
- resonant region, 4
- Schmidt decomposition, 11
- singular value decomposition, *see* SVD
- SPT, *see* symmetry protected topological phase
- subdiffusion, xi
- SVD, 11
- symmetry protected topological phase, x