



International School for Advanced Studies • Trieste, Italy
Academic year 2019–2020

Non-ergodic phenomena in many-body quantum systems

Author
TOMMASO PAROLINI

Supervisor
ANTONELLO SCARDICCHIO

Thesis submitted for the degree of
DOCTOR PHILOSOPHIAE
in the XI cycle of the
PhD course in Statistical Physics.

Abstract

The assumption of ergodicity is the cornerstone of conventional thermodynamics, connecting the equilibrium properties of macroscopic systems to the chaotic nature of the underlying microscopic dynamics, which eventuates in thermalization and the scrambling of information contained in any generic initial condition. The modern understanding of ergodicity in a quantum mechanical framework is encapsulated in the so-called eigenstate thermalization hypothesis, which asserts that thermalization of an isolated quantum system is a manifestation of the random-like character of individual eigenstates in the bulk of the spectrum of the system's Hamiltonian. In this work, we consider two major exceptions to the rule of generic thermalization in interacting many-body quantum systems: many-body localization, and quantum spin glasses.

In the first part, we debate the possibility of localization in a system endowed with a non-Abelian symmetry. We show that, in line with proposed theoretical arguments, such a system is probably delocalized in the thermodynamic limit, but the ergodization length scale is anomalously large, explaining the non-ergodic behavior observed in previous experimental and numerical works. A crucial feature of this system is the quasi-tensor-network nature of its eigenstates, which is dictated by the presence of nontrivial symmetry multiplets. As a consequence, ergodicity may only be restored by extensively large cascades of resonating spins, explaining the system's resistance to delocalization.

In the second part, we study the effects of non-ergodic behavior in glassy systems in relation to the possibility of speeding up classical algorithms via quantum resources, namely tunneling across tall free energy barriers. First, we define a pseudo-tunneling event in classical diffusion Monte Carlo (DMC) and characterize the corresponding tunneling rate. Our findings suggest that DMC is very efficient at tunneling in stoquastic problems even in the presence of frustrated couplings, asymptotically outperforming incoherent quantum tunneling. We also analyze in detail the impact of importance sampling, finding that it does not alter the scaling. Next, we study the so-called population transfer (PT) algorithm applied to the problem of energy matching in combinatorial problems. After summarizing some known results on a simpler model, we take the quantum random energy model as a testbed for a thorough, model-agnostic numerical characterization of the algorithm, including parameter setting and quality assessment. From the accessible system sizes, we observe no meaningful asymptotic speedup, but argue in favor of a better performance in more realistic energy landscapes.

Contents

Preface	xi
Notation	xiii
1 Quantum Dynamics and Quantum Algorithms	1
1.1 The ergodic hypothesis	1
1.1.1 Classical ergodicity	1
1.1.2 Quantum ergodicity	4
1.1.3 Ergodicity breaking	9
1.2 A brief history of spin glasses	13
1.3 Computing with quantum mechanics	20
1.3.1 A foreword on quantum computation	20
1.3.2 Quantum annealing	22
2 Many-Body Localization and Symmetry	29
2.1 Localized quantum systems	29
2.1.1 Anderson localization	30
2.1.2 The effect of interactions	33
2.2 MBL detectors and LIOMs	34
2.2.1 Phenomenological characterizations of MBL	34
2.2.2 The LIOM theory of MBL	41
2.3 Isotropic disordered models	43
2.3.1 Motivation: the Hubbard model with local disorder	44
2.3.2 Low-temperature thermodynamics	50
2.3.3 SU(2) symmetry	52
2.3.4 Symmetry and localization	56
2.4 Case study: the random-bond Heisenberg model	59
2.4.1 Preliminary definitions and remarks	59

2.4.2	SDRG for excited states	62
2.4.3	Exact diagonalization	67
2.4.4	Resonance counting	78
2.4.5	Conclusions	88
3	Simulating Quantum Tunneling	91
3.1	Quantum annealing and energy barriers	91
3.2	Quantum Monte Carlo methods	92
3.2.1	Path-integral Monte Carlo	92
3.2.2	Diffusion Monte Carlo	96
3.3	Simulating tunneling events on classical machines	100
3.3.1	Previous results and motivation	101
3.3.2	DMC tunneling in a quartic double-well potential	103
3.3.3	Semiclassical theory of the DMC tunneling dynamics	107
3.3.4	The ferromagnetic quantum Ising chain	109
3.3.5	The shamrock model	113
3.3.6	Finite-population bias	116
3.3.7	Conclusions	117
4	Population Transfer	119
4.1	Energy matching and the population transfer algorithm	119
4.2	PT in the impurity band model	123
4.2.1	Model definition and motivation	123
4.2.2	Semiclassical treatment and perturbation theory	125
4.2.3	PT time and quantum speedup	128
4.3	PT in the random energy model	131
4.3.1	Model definition and previous results	132
4.3.2	Quench dynamics in the non-ergodic phase	134
4.3.3	PT quality: spread and spillage	141
4.3.4	Quantum advantage	144
4.4	Conclusions	149
	Closing remarks	153
	Acknowledgments	157
A	The Jordan–Wigner transformation	159
B	Dimension of spin multiplets	163

C	Statistical properties of fusion trees	169
C.1	Distribution of nearest-neighbor graph distances	169
C.2	Size of the block spins	171
D	Correlated cost functions	173
E	Uniformity in Hamming space	181
	Bibliography	185

*To the beloved memory
of my grandfather.*

Preface

What is thermodynamic equilibrium? An answer echoes in the words of Feynman's: it is "fast things" being far in the past and "slow things" far in the future. During this Goldilocks period in the history of a system, macroscopic properties do not change in any appreciable manner and they can be ascribed a well-defined, stable value. Even from such an informal definition, the notion of equilibrium can be seen to crucially depend on the choice of observational time scale over which the system is considered, as we must specify what we deem "slow" and what "fast" for our current purposes.

Classical thermodynamics in its entirety aims at describing equilibrium states and the processes connecting them to one another. There is a good reason for this: they are far simpler than the alternative. As a matter of fact, the theory ultimately owes its success to the realization that systems at equilibrium require very little information to characterize to a perfectly satisfactory degree, reducing the exorbitant dimensionality of the generic, out-of-equilibrium many-body problem to a handful of state variables such as temperature, volume and pressure. It is worth noting that the very fact that equilibrium systems exist *at all* is a delightful property of our universe, no ascription of which to an anthropic principle — the universe *ought* to appear equilibrated on human scales if it is to support human life — can make any less wondrous.

Once thermodynamics was "figured out", it was natural for scientists to start turning their attention into the more exciting realm of out-of-equilibrium physics. This interest was enhanced, in the condensed matter world, by the recent technological developments in experimental physics, which made it possible to simulate a number of many-body Hamiltonians in a scaled up and much better controlled environment. As conventional thermodynamics does not work for non-equilibrium systems, novel tools need to be found in this endeavor, which encompasses the bulk of present-day research in many-body physics. Today, several communities are working in this very direction at the forefront of a broad range of disciplines, including statistical and condensed matter physics, quantum computation and information, computer science, and many more.

In this thesis, we explore some of the most interesting phenomena that have been

recently observed out of thermodynamical equilibrium. In our selection of examples, we have strived to convey the wide-ranging nature of the topic by focusing on both the physical understanding of the systems considered as well as their relationship with ongoing research in quantum computing.

We have structured the content as follows: in Chapter 1, functioning as an introduction, we discuss the role of the ergodic hypothesis in statistical physics. Its violation, which can occur in several different ways, is the *raison d'être* of this thesis. We then present a brief history of classical spin glasses, which constitute the first and most venerable class of such theory-defying systems, and for which a good deal of understanding has been achieved through the work of many starting from the late 1950s. Finally, we take the opportunity to introduce the concept of quantum-enhanced optimization through quantum annealing, which finds in spin glasses a natural class of target problems where to hopefully outspeed classical computers.

In Chapter 2 we present the concept of many-body localization, which is a generic mechanism of strong violation of ergodicity in a large class of interacting systems. In our current understanding, this violation stems from the appearance of an extensive set of local constants of the motion as a result of disorder-induced quantum dephasing. We discuss how disordered systems with non-Abelian symmetry groups are not expected to undergo many-body localization, and survey a paradigmatic case. We find that although symmetry does appear to protect the system from full localization, non-ergodic phenomena are still observed and expected to persist up to very large sizes, suggesting a new kind of non-ergodic phase.

In Chapter 3 we aim to compare the expected performance of quantum annealers against classical optimization in a “tunneling competition”. Quantum tunneling is believed to be the empowering force of quantum annealers in finding low-lying states of glassy systems. After an introduction to the path-integral and diffusion Monte Carlo algorithms, we make use of the latter to classically simulate tunneling phenomena both in a one-dimensional toy model, where analytical considerations are easier, and in many-body spin problems. We find a generic scaling of the “classical tunneling time” which outperforms incoherent quantum tunneling, and we comment on the consequences.

In Chapter 4 we again aim to compare classical and quantum algorithms, but with a focus on energy matching in “golf-course” energy landscapes, where local-search heuristics cannot generally outperform random or brute-force search. We numerically study the so-called population transfer algorithm applied to a toy problem with uncorrelated random energy levels. Although we find population transfer to outspeed random search in an oracular sense, we find no evidence of quantum speedup when accounting for the time required to implement an oracle call, at variance with previous results on a simpler toy model (which we present at the start of the chapter).

This thesis is based on Refs. [213] (Ch. 2), [130,202] (Ch. 3) and [203] (Ch. 4).

Notation

Numbers and units

We always adopt natural units in our equation, so $\hbar = k_B = c = 1$. When quoting experimental figures and similar, we often make use of SI units (m, s, kg, *etc.*) instead.

Numerical uncertainty is denoted by round brackets after the last significant digits, *e.g.* 1.40(13) means 1.40 ± 0.13 .

Typeface

Variables, parameters and generic constants are denoted by italic letters or Greek symbols, *e.g.* x, α, c_1 . We reserve roman type for the Euler and imaginary number, e and i respectively.

Vectors (and covectors) tend to be denoted by normal italic variables; so, for instance, both the one-dimensional and d -dimensional Lebesgue measures are denoted by dx rather than $dx^1 \cdots dx^d$, unless it is convenient to draw the distinction. However, we sometimes reserve boldface type for vectors in the Euclidean space, including operators which transform under the rotation group, such as $\mathbf{s}_i = (s_i^x, s_i^y, s_i^z)$. This makes it simpler to separate the spatial and internal degrees of freedom without a proliferation of abstract indices.

Standard variable names

We try to avoid confusion between the Hilbert space dimension, the spatial dimensionality and the number of particles by consistently denoting the first one by D (or some variant thereof, such as $D_{L,S}$), the second one by d , and the third one by n (in the general setting) or L (in one dimension, when it coincides with the number of lattice sites).

Binary operators

\wedge is sometimes used to shorthand the binary minimum (meet), *e.g.* $x_a \wedge b = \min\{x_a, x_b\}$, unless it denotes the wedge product of forms (the reader should have no trouble distinguishing the two uses). Likewise, the join \vee can be used for the binary maximum. This notation, which is not quite standard, is always locally explained.

We decided to eschew the \equiv or $:=$ notation for definitions, limiting ourselves to regular equalities instead. We use \approx for approximate numerical equality and \sim for asymptotic (up to prefactors) or order-of-magnitude equality.

Brackets

Angular brackets $\langle \cdot \rangle$ always denote an *ensemble average* of some kind. Unfortunately — but inevitably in quantum statistical mechanics — the exact meaning is context-dependent. Usually they refer to a *thermal average*, which in both classical and quantum mechanics means averaging over the Gibbs distribution. More generally, in a quantum setting they can denote the expectation of an operator on any (possibly nonthermal) state ρ , namely $\langle O \rangle = \text{Tr}[\rho O]$. We also extend the notation to non-observable quantities (such as the r -parameter or the entanglement entropy) whose mean is taken over some set of states. The precise meaning will be clarified in context. In Dirac's bra-ket notation, $\langle \cdot | \cdot \rangle$ denotes instead the standard inner product and there can be no confusion.

Double brackets $\llbracket \cdot \rrbracket$ are reserved for the *disorder average*, which is frequently used in the context of glassy physics and MBL. We also adopt the double average notation $\langle\langle \cdot \rangle\rangle = \llbracket \langle \cdot \rangle \rrbracket$.

Functions

A function is said to belong to $C^n(\mathcal{X})$ if it is n times continuously differentiable in \mathcal{X} .

We employ standard big-O notation:

$$\begin{aligned} f(x) \in O(g(x)) & \text{ if } f(x) \leq cg(x) \text{ for some } c > 0 \text{ in a neighborhood of } x = x_0; \\ f(x) \in \Theta(g(x)) & \text{ if } f(x) \in O(g(x)) \text{ and } g(x) \in O(f(x)); \\ f(x) \in o(g(x)) & \text{ if } f(x)/g(x) \rightarrow 0 \text{ as } x \rightarrow x_0; \\ f(x) \in \omega(g(x)) & \text{ if } g(x) \in o(f(x)); \end{aligned}$$

The limit point is usually $x_0 = \infty$ or $x_0 = 0$, and always specified when not obvious.

All logarithms whose base is left unspecified are taken in the natural base.

Chapter 1

Quantum Dynamics and Quantum Algorithms

1.1 The ergodic hypothesis

It is unfortunate that present prejudice is towards ergodic behaviour, so that its breakdown seems notable; ergodicity is really much more remarkable than its failure.

R. G. Palmer [198]

1.1.1 Classical ergodicity

The zeroth law of thermodynamics asserts that two systems simultaneously in thermal equilibrium with a third one must also be in thermal equilibrium with each other. The notion of thermal equilibrium, in turn, is defined from the empirical conditions of absence of heat flows between the systems. This allows one to define a temperature state function T , which labels the cosets of the “thermal equilibrium” equivalence relation and whose gradient causes heat flows between proximal regions, *e.g.* via Fourier’s law

$$\frac{d^2Q}{dt d\sigma} = -k\nabla T, \quad (1.1)$$

where k is a material-dependent heat conductivity and the left hand side represents the heat exchange rate per unit surface. Plausibly, and in accord with our daily-life observations, one then expects two systems at different temperatures to eventually reach thermal consensus when placed in contact: a process called *thermalization*.

The first attempt at a microscopic explanation of this process was the work of Boltzmann's [51], building on previous results by Maxwell [171–173]. Boltzmann considered an ideal gas and supposed that collisions between two particles happened with zero correlation between the particles' kinetic energies and positions (*Stoßzahlansatz*). With these assumptions, he could prove that a certain state function, later known as the H function, was nonincreasing in time, and in fact strictly decreasing until the attainment of its minimum value. This value corresponds to the ideal gas settling on a Maxwell–Boltzmann (thermal) distribution of particle velocities, whose temperature is fully dictated by the initial value of the energy.

Boltzmann was aware that neither his Ansatz nor the consequent H -theorem were tenable in the strictest sense, for instance violating classic recurrence theorems, but it could be argued (cf. [90]) that small deviations from it were so inconsequential, and large deviations so rare, that they could not harm the general idea of the system eventually reaching thermal equilibrium in every practical sense of the word.¹

In modern notation, we know that a classical system can be quite generally parametrized by a set of generalized coordinates and momenta, $x = (q, p) \in \Omega \subset \mathbb{R}^{2s}$,² living in a symplectic manifold called the *phase space*. Systems are then described in terms of *macrostates*, *i.e.* sets in the phase space characterized by specific values of a fixed collection of macroscopic properties (*e.g.* internal energy, volume, temperature, *etc.*). A macrostate is conveniently formalized as an *ensemble*, namely a probability distribution $\rho : \Omega \rightarrow \mathbb{R}_0^+$ with

$$\int_{\Omega} \rho(q, p) dq \wedge dp = 1. \quad (1.2)$$

The ensemble formalism provides us with a natural way of averaging microscopic functions f over the phase space:

$$\langle f \rangle = \int_{\Omega} \rho(q, p) f(q, p) dq \wedge dp. \quad (1.3)$$

The crucial link between thermodynamics and mechanics consists in interpreting each macroscopic property F_{obs} of a system as the ensemble average of some appropriate microscopic function f defined on Ω ,

$$F_{\text{obs}} = \langle f \rangle. \quad (1.4)$$

This is achieved in two steps.

First, notice that, from a mechanical standpoint, the system under consideration must be described by a particular instantaneous microstate $(q(t), p(t))$ whose law of motion in

¹This counter-objection was in a sense put on firmer grounds with the development of the information-theoretical understanding of entropy, which was also pioneered by Boltzmann through his most famous formula, $S = k_B \log |\Omega|$ for a system with $|\Omega|$ equiprobable microstates.

²Here s is the number of coordinates, *e.g.* $s = dn$ for n free particles in d dimensions.

the phase space is dictated by the Hamilton equations

$$\begin{cases} \frac{dq}{dt} = \nabla_p H, \\ \frac{dp}{dt} = -\nabla_q H \end{cases} \quad (1.5)$$

with appropriate initial conditions. It is important to point out that while the ensemble description formalizes our inability to exactly pinpoint the particular trajectory $(q(t), p(t))$, at least in classical mechanics we do admit that such a trajectory exists. As a consequence, whatever macroscopic property must simply reflect the value of some corresponding microscopic function $f = f(q, p)$, averaged along the trajectory over the observation time:

$$F_{\text{obs}} = \frac{1}{\tau_o} \int_{t_0}^{t_0 + \tau_o} f(q(t), p(t)) dt. \quad (1.6)$$

Here the experiment is defined to start at t_0 and end at $t_0 + \tau_o$, and if the measurement is to be well-defined, we should as well assume that the system is at equilibrium, that is to say, any time scale characterizing a significant variation of $f(q(t), p(t))$ must be either much smaller or much larger than the experiment's duration τ_o . If that is the case, we may as well drop all the extremely slow degrees of freedom from the Hamiltonian description and take the $\tau_o \rightarrow \infty$ limit in Eq. (1.6).

The second step is more delicate. While Eq. (1.6) provides an explicit connection between theory and experiments, it is utterly useless from a practical standpoint, as it still presupposes a knowledge of the exact microstate which is neither attainable nor desirable for a many-body system. What we would rather have instead, is to link the value of F_{obs} to the ensemble representation introduced above, or in other words, to switch from a mechanic to a thermodynamic description.

The missing link is provided by the *ergodic hypothesis*. In its basic form, it states that infinite-time averages can be replaced by ensemble averages:

Definition 1.1.1. (Ergodic hypothesis) Given a Hamiltonian system, for any time t_0 , any initial condition $x(t_0) = x_0$ and any measurable function $f = f(q, p)$, one has

$$\lim_{\tau_o \rightarrow \infty} \frac{1}{\tau_o} \int_{t_0}^{t_0 + \tau_o} f(q(t), p(t)) dt = \langle f \rangle \quad (1.7)$$

for some ensemble $\rho = \rho(q, p)$ which is only determined by x_0 . Here $\langle f \rangle$ is defined by Eq. (1.3) and the trajectory $(q(t), p(t))$ is determined by integrating Eqs. (1.5) from the initial condition.

The special ensemble verifying Eq. (1.7) should obviously describe the system at equilibrium. Its explicit form can be derived from a reasonable assumption of unbiasedness.

Indeed, given some macroscopic information on the system and two microstates compatible with such information, one would be hard pressed to justify assigning different likelihoods to the two. The only distribution that agrees with this principle is the (normalized) indicator function of the set of compatible microstates.

When the only available information is the system's energy, the corresponding pdf is called the *microcanonical ensemble* ρ_{mc} . It assigns uniform probability to the “energy shell”

$$\Omega(E, \delta E) = \{x \in \Omega \mid E - \delta E/2 \leq H(x) \leq E + \delta E/2\} \quad (1.8)$$

and zero elsewhere. The thickness δE is to be taken infinitesimally small, $\delta E/E \ll 1$.

The term *ergodic*, introduced by Boltzmann in the late 1870s, is nowadays used in the theory of dynamical systems to describe a measure-preserving flow $\phi_t : \Omega \rightarrow \Omega$ on some probability space $(\Omega, \mathcal{O}, \mathbb{P})$ with the property that, for any $t \in \mathbb{R}$ and any $X \in \mathcal{O}$,

$$\phi_t^{-1}(X) \subset X \implies \mathbb{P}(X) \in \{0, 1\}. \quad (1.9)$$

This means that no set other than the whole space Ω , up to null-sets, is left invariant by ϕ_t , and as a corollary, the flow causes \mathbb{P} -almost every point to visit every single neighborhood of the space. Closed orbits, including trivial ones, must add up to a zero-measure set. In physical terms, there is no nontrivial constant of the motion other than energy.

In the Gibbsian formulation of statistical mechanics, ergodicity is assumed of the Hamiltonian flow on the phase space, or rather on its shells $\Omega(E, \delta E)$, with respect to the measure induced by the microcanonical ensemble.

From this assumption³ it is possible to derive Eq. (1.7) [44, 45, 259] (see also Refs. [188, 246] for a historical overview and in-depth discussion), and indeed one may take it as a slightly stronger definition of the ergodic hypothesis.

1.1.2 Quantum ergodicity

The task of reformulating statistical mechanics in the quantum theory was commenced very early in the history of quantum mechanics. The most important early attempt at a rigorous formulation of a quantum concept of ergodicity must be recognized in von Neumann's work in the late 1920s [258]. He defined a microcanonical ensemble by subdividing the spectrum of the Hamiltonian operator, assumed to be discrete, into “slices” (shells) which are macroscopically thin, $\delta E \ll E$, but still contain exponentially many states.

³Actually, the assumption may be relaxed to that of *metric transitivity*, i.e. the flow should eventually cause any non-null measurable set of the phase space to overlap with any other. In formulas,

$$\forall A, B : \mathbb{P}(A), \mathbb{P}(B) > 0, \exists t : \mathbb{P}(\phi_t^{-1}(A) \cap B) > 0. \quad (1.10)$$

Eq. (1.10) implies ergodicity in the sense of Eq. (1.9), or else we could find two distinct ϕ_t -invariant sets A and A^c with nonzero measure, leading to a contradiction: $0 < \mathbb{P}(\phi_t^{-1}(A) \cap A^c) = \mathbb{P}(A \cap A^c) = 0$. The opposite implication is not true in general, but it does hold for stationary flows.

Then, the Hilbert space \mathcal{H} associated to each shell is in turn partitioned into macrostates, namely, orthogonal subspaces corresponding to the different values we can measure for a prescribed set of macroscopic observables:

$$\mathcal{H} = \bigoplus_v \mathcal{H}_v, \quad (1.11)$$

with \mathcal{H}_v the common eigenspace of a set of commuting observables M_1, \dots, M_m relative to the respective eigenvalues m_1, \dots, m_m . If the dimensions of \mathcal{H} and \mathcal{H}_v are D and D_v , respectively, the microcanonical probability associated to the macrostate \mathcal{H}_v is just D_v/D .

Von Neumann proved the following result, which we state non-rigorously: for most Hamiltonians and most macro-partitions $\{\mathcal{H}_v\}_v$, every initial state $|\psi_0\rangle$ reproduces the microcanonical distribution for most of the time during its evolution, in the sense that

$$\mathbb{E} \left[\|P_v \psi(t)\|^2 \right] = \frac{D_v}{D}, \quad (1.12)$$

$$\text{Var} \left[\|P_v \psi(t)\|^2 \right] < \frac{1}{D_v} \left(\frac{D_v}{D} \right)^2 \quad (1.13)$$

for an arbitrarily large fraction of times t , where $|\psi(t)\rangle = e^{-iHt} |\psi_0\rangle$ and P_v is the projector onto \mathcal{H}_v [114].

This theorem, which von Neumann called a “quantum ergodic theorem” (QET), is clearly reminiscent of the ergodic hypothesis in classical mechanics. However, notice a crucial difference. If we were to strictly parallel the classical formalism, the way to go about defining quantum ergodicity would be to require the long-time average to equal the ensemble average, as in Eq. (1.7). In a quantum context, the appropriate ensemble for long-time averages is the *diagonal ensemble*: if $|\psi_0\rangle = \sum_a c_a |E_a\rangle$ is a decomposition of the initial state into eigenstates of H with E_a within the chosen energy shell (*i.e.* we are considering an initial state with a well-defined energy density), then

$$|\psi(t)\rangle \langle \psi(t)| = \sum_{ab} c_a c_b^* e^{-i(E_a - E_b)t} |E_a\rangle \langle E_b|, \quad (1.14)$$

which implies

$$\begin{aligned} \int_0^{\mathcal{T}} \frac{dt}{\mathcal{T}} |\psi(t)\rangle \langle \psi(t)| &= \sum_{ab} c_a c_b^* \int_0^{\mathcal{T}} \frac{dt}{\mathcal{T}} e^{-i(E_a - E_b)t} |E_a\rangle \langle E_b| \\ &\rightarrow \sum_a |c_a|^2 |E_a\rangle \langle E_a| \quad \text{as } \mathcal{T} \rightarrow \infty, \end{aligned} \quad (1.15)$$

where we used the assumption of nondegenerate level spacings to turn the time average of the fluctuating phase factor into a δ_{ab} . The convergence must be intended in a weak sense,

i.e., identity under tracing. On the other hand, the microcanonical ensemble corresponds to the uniform distribution on the shell, namely

$$\rho_{\text{mc}} = \frac{1}{D} \sum_a |E_a\rangle \langle E_a|, \quad (1.16)$$

with the sum running over an eigenbasis basis of \mathcal{H} . In order for Eqs. (1.15) and (1.16) to coincide, it must be the case that $|c_a|^2 = 1/D$ for all a , an extremely nongeneric initial condition.

We see therefore that the strict quantum analogue of the ergodic hypothesis must be rejected. Indeed, von Neumann's QET suggests that the time/ensemble average equivalence ought to be pursued not on the level of probability distributions, but only on a coarser scale — that of macrostates. More precisely, we can say that a microcanonical description is appropriate for the system provided ρ_{mc} and the diagonal ensemble assign the same average values to all macroscopical observables, that is to say, to all macrostate projectors P_ν :

$$\sum_a |c_a|^2 \langle E_a | P_\nu | E_a \rangle \approx \frac{D_\nu}{D} \quad \text{for all } \nu, \quad (1.17)$$

with the \approx sign admitting fluctuations of the order of Eq. (1.13). Von Neumann's QET ensures that for typical Hamiltonians and macro-partitions, condition (1.17) holds for *all* choices $\{c_a\}_a$ of initial conditions. Moreover, it holds not only for the diagonal (*i.e.* long-time average) ensemble, but even for the instantaneous ensemble $|\psi(t)\rangle \langle \psi(t)|$, for most times t .

Von Neumann's result went largely unnoticed or misattributed for a long time, and after the 1950s, it started being misquoted and unjustly accused of triviality as a consequence (see a discussion in Ref. [114]). In the meantime, thermalization in quantum systems was typically modeled by the use of the mixed-state formalism and the idea of coupling to an external bath. It was not until the 1990s that the idea of diagonal/microcanonical equivalence resurfaced in the foundational works of Deutsch [84] and Srednicki [239].

As Deutsch pointed out, the heat-bath stratagem, while practically useful, is conceptually unsatisfactory as it still leaves the question unanswered on how the global system (*i.e.*, including the heat bath in the density matrix description) can achieve thermalization. Deutsch considered instead *isolated* systems, showing how integrable Hamiltonians could generically be made ergodic (in the sense of macro-equivalence of ρ_{mc} and the diagonal ensemble) by the addition of uncorrelated Gaussian noise (a GOE matrix; cf. Section 2.2) with intensity larger than the level spacing, but still arbitrarily small in the thermodynamic limit. The effect of the noise, Deutsch argued, is to create huge superpositions of states such that when averaging a macro-observable over an eigenstate, an exponential number D of similarly-distributed terms is summed together and quantum fluctuations are suppressed as a consequence as $\sim 1/\sqrt{D}$. This makes the diagonal ensemble equivalent to the microcanonical one, *i.e.* the system becomes ergodic [84].

Srednicki, attacking the problem from the perspective of chaos theory, used a similar strategy by positing Berry's random-wave conjecture [41], where eigenstate wave functions essentially behave as uncorrelated Gaussian vectors. Working in momentum space for a gas of hard spheres, he proved that with this assumption the single-particle momentum distribution for an eigenstates sharply tends to the thermal one, implying thermalization of generic initial conditions [239].

The modern understanding of *eigenstate thermalization*, which has taken the place of ergodicity in quantum statistical mechanics, can be formulated as follows [215,240,241]:

Definition 1.1.2. (Eigenstate thermalization hypothesis) Let H be a Hamiltonian on n particles with nondegenerate eigenvalues $\{E_a\}_a$ and respective eigenstates $\{|E_a\rangle\}_a$. H is said to obey the *eigenstate thermalization hypothesis* (ETH) in the interval $\mathcal{I} \subset \mathbb{R}$ if for every $E_a, E_b \in \mathcal{I}$ and every local⁴ observable O ,

$$\langle E_a | O | E_b \rangle = \delta_{ab} O(E) + e^{-S(E)/2} f(E, \omega) R_{ab} \quad (1.18)$$

where $E = \frac{E_a + E_b}{2}$, $\omega = E_a - E_b$, $O(E)$ is continuous in its argument, $f(x, y)$ is continuous in both arguments and $O(n^0)$, $S(E)$ is the microcanonical entropy (natural logarithm of the density of states) and R_{ab} is a random variable with a Gaussian distribution of zero mean and unit variance.

There is a lot of information packed in Eq. (1.18), so let us go through it. First off, notice that the ETH implies the diagonal/microcanonical ensemble identity:

$$\sum_a |c_a|^2 \langle E_a | O | E_a \rangle = \sum_a |c_a|^2 O(E_a) \approx O(E), \quad (1.19)$$

where we took a state $\psi = \sum_a c_a |E_a\rangle$ living in the energy shell centered at E , *i.e.* $c_a = 0$ whenever $E_a \notin \Omega(E, \delta E)$. This allows us to take $O(E)$ out of the sum, which then disappears: $\sum_a |c_a|^2 = 1$. Incidentally, Eq. (1.19) shows that $O(E)$ in Eq. (1.18) must in fact be the microcanonical average of O :

$$O(E) = \text{Tr}[\rho_{\text{mc}}(E)O]. \quad (1.20)$$

The second term at the right-hand side of Eq. (1.18) can be understood by considering the time-averaged deviation of the instantaneous quantum expectation of O from the microcanonical prediction: if we denote the long-time average with an overline,

$$\overline{O} = \lim_{\mathcal{T} \rightarrow \infty} \int_0^{\mathcal{T}} \frac{dt}{\mathcal{T}} O(t) \quad (1.21)$$

⁴The precise meaning of "local" depends on the underlying geometry. For a lattice Hamiltonian, we can simply take O to be supported on a finite, $O(n^0)$ number of sites. In continuous space, the essential support of O should be bounded in diameter. In both cases, we could relax the constraint somewhat and turn a blind eye to "exponentially-decaying tails".

with $O(t) = \langle \psi(t) | O | \psi(t) \rangle$, we have

$$\overline{(O(t) - \bar{O})^2} = \sum_{ab} |c_a c_b f(E, \omega) R_{ab}|^2 e^{-S(E)}, \quad (1.22)$$

which is $O(e^{-S(E)})$ for a state living in the shell $\Omega(E, \delta E)$. Notice that for all energies in the bulk of the spectrum we have $S(E) \sim n$, so the fluctuations away from the microcanonical value $\bar{O} = O(E)$ are exponentially suppressed with the system size. This results in an approximately one-valued relation between the matrix elements and the energy

It is worth stressing that ETH contains more than just the time/ensemble equivalence. Indeed, while the latter is only concerned with averages of macroscopic observables, as in Von Neumann's QET (cf. Eq. (1.17)), the former is a statement about *all* bulk eigenstates of the system, *i.e.* a microscopical hypothesis. Remarkably, it asserts that all eigenstates in a sufficiently narrow energy window "look alike", in the sense that no local measurement can distinguish between them. This implies as well that a *single* eigenstate contains all the information relevant for thermodynamics, as the thermal value of any local observable is equivalent to the expectation value of that observable in any one eigenstate with the appropriate energy density. This observation highlights a profound difference between the concept of thermalization in classical versus quantum mechanics. In the words of Ref. [215], according to ETH "every eigenstate of the Hamiltonian always implicitly contains a thermal state. The coherence between the eigenstates initially hides it, but time dynamics reveals it through dephasing."

We shall make use of the ETH as a stringent test of ergodicity in Section 2.4.3; in particular, see Fig. 2.13 for a concrete manifestation of the ETH (as well as its violation) for two local observables probed in the eigenbasis of a spin Hamiltonian.

To conclude, we provide a formal definition of what it actually means for a quantum system to thermalize. We already mentioned how the old approach to quantum statistical mechanics, which defined thermalization by introducing an external heat reservoir, cannot fully appease the theoretician. The main issue is that a quantum system considered in its entirety should be regarded as *isolated*, as external couplings can always be incorporated (in principle) into the system's Hamiltonian, until everything is accounted for. But an isolated system evolves unitarily, and if we take it to start from an initial pure state (as it is always possible by enlarging the Hilbert space sufficiently), it will never become mixed.

The solution to this mathematical paradox is physical in nature. We do not need to require the full system to be described by a thermal ensemble in order to *observe* thermalization. We only need it to thermalize *locally*. This is expressed in the following manner [192]:

Definition 1.1.3. (Local thermalization) A state $|\psi_0\rangle$ undergoing time evolution under the Hamiltonian H , $|\psi(t)\rangle = e^{-iHt} |\psi_0\rangle$, is said to (*locally*) *thermalize* if for any bounded spatial region A , the reduced density matrix

$$\rho_A(t) = \text{Tr}_{A^c} [|\psi(t)\rangle \langle \psi(t)|] \quad (1.23)$$

behaves, in the large-time and large-system size limit, as though it were traced from the Gibbs distribution. In other words, if

$$\lim_{t,n \rightarrow \infty} \rho_A(t) = \lim_{t,n \rightarrow \infty} \text{Tr}_{A^c} \left[\frac{e^{-\beta H}}{Z} \right] \quad (1.24)$$

for all bounded⁵ sets A , where the limit is taken with t/n fixed and the $=$ sign denotes identity under tracing.

In Eq. (1.24), $Z = \text{Tr}[e^{-\beta H}]$ is the partition function and β is an A -independent parameter fixed by the conservation of energy:

$$\lim_{t,n \rightarrow \infty} \text{Tr} \left[\frac{e^{-\beta H}}{Z} H \right] = \langle \psi_0 | H | \psi_0 \rangle. \quad (1.25)$$

Notice that we are not requiring $\lim_{t,n \rightarrow \infty} \rho(t)$ to be a thermal state — which would be impossible for a pure $\rho(0)$. Indeed, the full density matrix does not even admit such a limit in general. However, considering local restrictions enables us to make the notion of infinite-time limit well-defined in the thermodynamic limit.⁶

Local thermalization is best pictured as the proper formalization of the old coupling-to-a-heat-bath prescription. Rather than introducing an *ad hoc* external reservoir for the local system to equilibrate with, the system itself acts as a reservoir for its own subsystems. It is a nice and desirable property that in order for this to occur, a *global* temperature must exist, determined by Eq. (1.25). Indeed, if two bounded sets A_1 and A_2 obeyed Eq. (1.24) with different β 's, their union $A_1 \cup A_2$ (still a bounded set) could not with any. At the end of the day, the reservoir formalism of quantum statistical mechanics is thus rescued, and the theoretician can rest easy.

1.1.3 Ergodicity breaking

So far, we have taken for granted that systems, whether classical or quantum, should generically thermalize, a notion which we identified with the existence of an equilibrium ensemble which attracts almost all initial conditions with a given average energy. To make everything well-defined — let it not be forgotten — the thermodynamic limit (or at least large enough system sizes) must be considered.

It is the main focus of this thesis to show how this assumption can be violated in certain situations. In particular, in Chapter 2 we will go over the concept of Anderson

⁵We can relax the boundedness requirement for A so long as the ratio of $|A|/|A^c|$ goes to zero in the thermodynamic limit.

⁶The simultaneous t and n limit prescription is necessary because in finite systems equilibration cannot strictly occur due to recurrences (similar to the phenomenon of Poincaré recurrence in classical dynamical systems), so $t \rightarrow \infty$ followed by $n \rightarrow \infty$ is never well defined. Likewise, $n \rightarrow \infty$ followed by $t \rightarrow \infty$ does not allot enough time for the system to possibly mix information to a satisfactory degree.

localization and many-body localization, where a quantum particle or a multitude of them is prevented from diffusing by the presence of impurities. In this section, we mention other mechanisms whereby a system can fail to achieve thermal equilibrium for an exponentially (in the system size) long time. We essentially follow a presentation of the subject matter by Palmer [198].

The simplest ergodicity-breaking mechanism we can mention is *spontaneous symmetry breaking* (SSB). In certain circumstances, a Hamiltonian with a certain global symmetry might have low-energy states which are not invariant under that symmetry. For instance, we can have a nontrivial orbit of the symmetry group containing multiple (possibly infinitely many) ground states. A broken symmetry can partition the configuration space into distinct, dynamically separate *components*. The prototypical example is the classical ferromagnetic Ising model in the low-temperature phase, whose two ground states, corresponding to the upward and downward totally-polarized configurations, break the \mathbb{Z}_2 symmetry of the Hamiltonian. They are dynamically separate in the sense that the free energy barrier ΔF between them (which can be seen as a double-well potential, cf. Section 3.3) goes to infinity in the large-size limit, preventing the system from switching between the two, as would be required by the ergodic hypothesis. Indeed, if the system is at equilibrium *within* component Γ at temperature $T = 1/\beta$, the rate of escape can be estimated at⁷

$$\frac{P_{\text{esc}}(t)}{t} \sim \frac{Z(\partial\Gamma)}{Z(\Gamma)} = e^{-\beta\Delta F}, \quad (1.26)$$

which means that on average $t_{\text{esc}} \sim e^{\beta\Delta F}$. The system remains trapped inside Γ if this time is huge relative to the observational time scale τ_o ,

$$t_{\text{esc}} \gg \tau_o, \quad (1.27)$$

or equivalently, if the probability of escaping within the observational time scale is very small, $P_{\text{esc}}(\tau_o) \sim \tau_o e^{-\beta\Delta F} = \tau_o/t_{\text{esc}} \ll 1$. Clearly, if $\Delta F \sim n$ any fixed value for τ_o will indeed pale in comparison to $t_{\text{esc}} \in O(e^{cn})$ as soon as n is moderately large (see Fig. 1.1).

For this reason, equilibrium statistical mechanics fails to describe the actual properties of the system, for instance predicting zero magnetization for the Ising ground state instead of the physically correct $M \in \{\pm n\}$. This is equivalent to “getting the order wrong” when taking both the limits $n \rightarrow \infty$ (to be taken first) and $t \rightarrow \infty$ (afterwards) — a mistake which students are commonly warned against in any statistical physics course. The phenomenon of trapping is also called *metastability* when t_{esc} and τ_o are of the same order of magnitude.

We can recognize from the Ising example two main features that allow us to talk about the “splitting into components” of the configuration space:

- a) the probability of escape from any component is “small” within the allotted observation time: $P_{\text{esc}}(\tau_o) \ll 1$;

⁷Obviously, this rate is only well-defined in the $P_{\text{esc}}(t) \ll 1$ region, where $P_{\text{esc}}(t)$ increases roughly linearly with time. Eventually, $P_{\text{esc}}(t)$ will tend to 1 in the (unphysical) $t \rightarrow \infty$ limit for any fixed n .

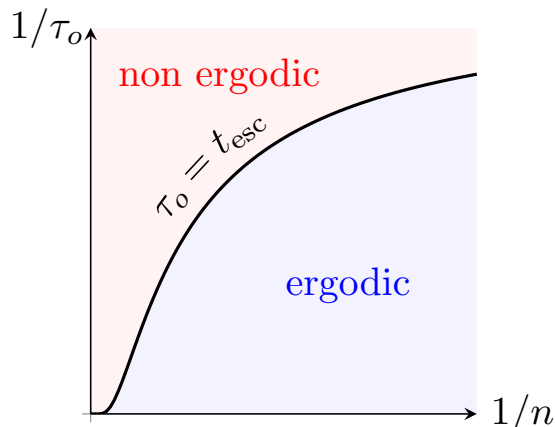


Figure 1.1: The definition of ergodicity depends on the choice of observational time scale τ_0 . When this is large compared to the typical time t_{esc} taken by the system to escape a confined region of the configuration space, the system is ergodic. A system with extensive free energy barriers between its components has $t_{\text{esc}} \in O(e^{cn})$, and as a result, it will always be non-ergodic at large enough sizes, for any fixed τ_0 .

b) the dynamics *within* each component is ergodic.

This prescription is general enough to not only capture the case of SSB, but also systems where no symmetry breaking occurs. Palmer gives the example of the H_2 molecule, whose ortho- and para- isomers are not related by a global symmetry transformation, but still define long-lived⁸ components of the configuration space.

Still, one may consider the above examples to be sort of trivial. Indeed, although the strict application of the canonical prescriptions of statistical mechanics fail to provide a proper description of the systems, it is not too hard to recognize where their limitations arise and, more importantly, how to fix them. For instance, in the Ising case one may simply add an arbitrarily small magnetic field to select either one of the ground states. This perturbation is energetically insignificant in the thermodynamic limit, but has the effect of properly accounting for the trapping mechanism and hence giving the right results using the canonical formalism.

In general, there are two ways of dealing with a restricted ensemble: the first one is to only trace on the degrees of freedom pertaining to the component of interest when calculating the partition function or other equilibrium quantities (the remaining degrees

⁸"In the low-density gas phase [ortho–para] conversion is extremely slow and accurate values are not known as measurements are usually perturbed by walls or magnetic contaminants such as O_2 . Out-of-equilibrium samples of H_2 at NTP [normal temperature and pressure] will typically convert a percent or two in a week's time; properly stored samples of D_2 at NTP have been known to show little change in the period of one year." [232]

of freedom being treated as nondynamical parameters); the Born–Oppenheimer approximation can be seen as a famous example. The second one is to modify the Hamiltonian itself, which includes not only the addition of artificial external fields as in the example above, but also other prescriptions such as imposing fixed boundary conditions or the mean-field approximation.

Far less trivial instances of ergodicity breaking are however possible. First off, even in the case of a few-component partition, there is no general rule on how to properly restrict the ensemble. In many occasions, heuristics or *a posteriori* justifications are used to guide the choice. Indeed, an ensemble restriction conceptually requires the injection of additional information into the canonical formalism, which by itself is equivalent to the Gibbs prescription of using the *least possible biased estimator* of a prior compatible with energy conservation (cf. the comment after Definition 1.1.1). This additional information has to be provided by the person operating the restriction. Again, in the Ising case it is not hard to see that *e.g.* the addition of an external field is equivalent to minimizing bias subject to the known average value of both energy *and* magnetization (*i.e.* the magnetic field is “conjugate” to M the same way the inverse temperature is to E). Lacking a clear-cut or simple enough order parameter such as M , the restriction process loses its transparency.

A particularly unwieldy scenario occurs when the number of components grows out of control, *e.g.* exponentially with n . In this case, an order parameter (or at least some parametrization of the component space) may still be possible to define, but one cannot hope to be able to actually measure it. In such cases, the only sensible option is to treat the component label as a random variable, with some given “disorder distribution”. The physical properties of the system are then obtained by a *double averaging* process: both over the restricted ensemble and over the disorder. This is the typical situation when treating *glasses*, and we shall accordingly encounter this “two-level statistical mechanics” (in Palmer’s words) in the next section.

Glasses are, generally speaking, systems where the dynamics undergoes a dramatic slowdown despite the system not being in its thermally-equilibrated state quite yet. This effectively freezes the system out-of-equilibrium, up to extremely slow relaxation processes called *aging* (see *e.g.* [163]). We are going to focus our attention on *spin glasses*, where the frozen degrees of freedom are spins and the mechanism underpinning the dynamical arrest is *frustration*. Other kinds of glassy systems include structural glasses [183], and systems such as granular materials and colloids, which undergo a jamming transition at high densities highly reminiscent of the glass transition [46].

With glassy physics already presenting a remarkably complex phenomenology at the classical level, it should come as no surprise that introducing quantum effects can make things even more interesting (and less clear). One could for instance study the jamming transition [25] or the spin glass transition [189] in a quantum setting. In general, the interplay of quantum mechanics and glassy physics is a rather new topic in physics, and the subject of active research. We will shirk the endeavor of exploring this matter in more detail, and content ourselves with presenting an introduction to *classical* spin glasses with

a little historical context.

1.2 A brief history of spin glasses

Spin glasses, which have been deservedly described as “the most complex kind of condensed state encountered [...] in solid state physics” [105], made their appearance around the 1950s.⁹ In those days, multiple experimental groups were interested in probing the magnetic properties of metallic alloys in which a magnetic and a nonmagnetic component are melted together. This typically involves embedding transition metals or ions — whose electronic *d*-shell is only partially filled and strongly interplays with the substrate’s conduction electrons — into inert, noble metals. Important examples include copper, silver and gold doped with manganese, iron or chromium [27].

A relevant parameter is the concentration of magnetic impurities, namely their density relative to that of the underlying paramagnetic matrix. Different concentrations often result in very distinct physical properties. In the low-concentration regime (typically, less than 0.1% dilution) magnetic impurities are essentially isolated and their only effect on the sample is to scatter conduction electrons and thereby amplify the alloy’s electrical resistivity at small temperatures, a mechanism called *Kondo effect* [7, 148, 254].

However, at moderate concentrations (*e.g.* 0.2–14% for Cu–Mn alloys [256]) the dominant effect is the magnetic interaction between impurities. Interestingly, in most cases this interaction is not governed by direct exchange between magnetic dipoles, but is rather mediated by the band electron’s rearrangement as a consequence of the ion’s positive charge (screening). An apt description of the resulting effective interaction between impurities was first achieved by Ruderman and Kittel [221] and later perfected by Kasuya [141] and Yosida [266], leading to the standard form of the *RKKY exchange integral* [86]

$$J(r) \sim -\frac{\pi}{E_F}(3nJ_0)^2 \frac{\sin(2k_F r - 3\pi/2)}{(2k_F r)^3} \quad (1.28)$$

for two impurities located at (large) distance r at a concentration $1/n$, J_0 being the direct coupling of the atom to the conduction electrons and k_F the Fermi wavenumber.

When $1/2k_F$ is smaller than the typical distance a_0 between ions, the sign of the interaction alternates between positive and negative — a phenomenon closely resembling Friedel oscillations in a Fermi gas in the presence of charged impurities [109]. In both cases, the oscillating term comes from the sharpness of the Fermi surface and is therefore screened out when the temperature is increased [102].

The position of ions in dilute alloys is highly irregular and can be approximated as a random lattice. It is equivalent, but conceptually simpler, to treat the lattice as regular and randomize the interactions (sign and amplitude) instead. This results in a typical *frustrated*

⁹See Ref. [21] and the following columns in the series for a fascinating debriefing on spin glass history.

spin Hamiltonian of the form

$$H = \sum_{\langle ij \rangle} \sum_{a,b} J_{ij}^a s_i^a s_j^b + \sum_i \sum_a h_i^a s_i^a, \quad (1.29)$$

where $a, b \in \{x^1, \dots, x^d\}$ and s_i^a is the a -component of a classical spin variable located at site i of the lattice; the notation $\langle ij \rangle$ means that we only sum on neighboring sites in the interaction lattice (meaning we are cutting off the $\sim 1/r^3$ interaction, which after all we know to be screened at large distance by a $\sim e^{-\alpha Tr/k_F}$ term at temperature T).

The term *frustration* refers to the fact that a typical lattice site is interacting with other sites both ferro- and antiferromagnetically, meaning that a single spin flip induces an unpredictable change in the overall energy. More importantly, the minimization (or maximization) of the global energy cannot be performed by simply minimizing each spin's local energy, as would be the case for a purely ferromagnetic model; rather, any spin assignment will inevitably "appease" certain interactions and "frustrate" others. The central role of frustration in the physics of spin glasses was first pointed out by Toulouse [252, 255].

Once this general feature is acknowledged, it can be usefully abstracted out and recognized elsewhere. Situations where one is trying to minimize some target function with many degrees of freedom, and which does not admit a trivial factorization in terms of local optimizers, are ubiquitous not only in physics, but in computer science, biochemistry, logistics, economics, *etc.* It is the realm of *optimization problems*.

A useful analogy is provided by the game of Go. In Go, a player is trying to maximize his territory and minimize his opponent's by placing one stone at a time on the intersections of a grid, in alternating turns. When a "fight" arises in some area of the board, there exist theoretical sequences of moves called *jōseki* that aim at locally optimizing the resulting territory or influence. However, the board counts 19×19 intersections. Expert players know that limiting oneself to *jōseki* fights is not a recipe for victory, because once the board starts to get populated one has to account for the entirety of its structure, rather than individual elements, and it may be convenient to sacrifice a local advantage in exchange for a larger global benefit.

Likewise, frustrated Hamiltonians (or *cost functions*, as they are more generally called in optimization theory) cannot be minimized by means of local "jōseki". On the one hand, this may sound discouraging as in the case of large (or even moderate-size) optimization problems we are not afforded a bird's eye view of the "board" (configuration space), whose cardinality increases exponentially in the problem size. On the other hand, unlike in Go we almost never aim at perfect minimization, but content ourselves with approximate solutions. While these are still very hard to achieve, they can oftentimes be reached even by local algorithms.

Let us go back to our spin-glass Hamiltonian, Eq. (1.29). For definiteness, we specialize

it to the form

$$H = \sum_{\langle ij \rangle} J_{ij} s_i s_j, \quad (1.30)$$

which is called the *Edwards–Anderson model* [89].¹⁰ The underlying lattice is regular and d -dimensional.

How can we characterize the onset of a glassy phase in the model? A natural way is to consider the canonical ensemble at a certain temperature T and look for the presence of magnetic ordering in the system, which can be measured by the average autocorrelation function of the local magnetization:

$$q_{\text{dyn}} = \lim_{\mathcal{T} \rightarrow \infty} \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=0}^{n-1} \overline{\langle s_i(0) s_i(t) \rangle}, \quad (1.31)$$

with the overline denoting temporal average in the range $t \in [0, \mathcal{T}]$ as in Eq. (1.21), angular brackets the thermal average $\langle O(s) \rangle = Z^{-1} \sum_s e^{-H(s)/T} O(s)$ over the choice of initial configuration $s(0) = s$, and per the discussion of Section 1.1 the order of limits is fundamental in ensuring a nontrivial result. A nonvanishing value of the quantity q_{dyn} denotes that the spins “prefer” on average to align in a particular direction, signaling magnetic ordering for this choice of temperature and disorder $\{J_{ij}\}_{ij}$. If the system were ergodic, q_{dyn} could be replaced by its equilibrated value

$$q_{\text{EA}} = \frac{1}{n} \sum_{i=1}^n \langle s_i \rangle^2 \quad (1.32)$$

(with the thermodynamic limit taken last), which is always zero as a consequence of the $s_i \mapsto -s_i$ symmetry of Hamiltonian (1.30). This is exactly the same phenomenon as the vanishing of the magnetization-per-spin,

$$m = \frac{1}{n} \sum_{i=1}^n \langle s_i \rangle \quad (1.33)$$

in ferromagnetic systems when care is not taken to restrict the ensemble average to the appropriate component. A standard way to operate this restriction is to add a small, uniform magnetic field h to the system and take $h \rightarrow 0$ *after* all the other limits have been performed. This results in a magnetization which is zero in the paramagnetic phase and nonzero in the ferromagnetic one, making it a good order parameter.

Similar considerations apply to the case of spin glasses. At low temperature, the configuration space splits into a large number of Palmer components (traditionally called “pure

¹⁰Edwards and Anderson considered an isotropic model, but we can limit ourself to classical Ising spins without losing much.

states” in the spin glass literature, and “clusters” in optimization); in each pure state, the low-lying spin configurations (*i.e.* those significantly contributing to the low-temperature free energy) tend to have a “frozen core”, *i.e.* the value of s_i , for extensively many i 's, will be the same for most configurations [182], giving $\langle s_i \rangle^2 > 0$ *within the component*. In order to get a sensible result for q_{EA} , one must first restrict the thermal average to single components, and later average over all the components weighted according to their individual free energies.¹¹ In this circumstance, we still have $m = 0$ as the frozen values of $\langle s_i \rangle$ have random signs for different i 's, eventually summing up to zero (or rather, to $O(\sqrt{n})$, which is then dominated by the $1/n$ prefactor). However, with this prescription the quantity q_{EA} defined in Eq. (1.32), known as the *Edwards–Anderson order parameter*, will take a nonzero value, discriminating between the paramagnetic and spin-glass phases.

A physical interpretation of the EA order parameter is that it quantifies the reduction in the magnetic susceptibility of the material compared to the paramagnetic value [105]. Let

$$\chi_{ij} = \frac{\partial m_i}{\partial h_j} \quad (1.34)$$

be the susceptibility tensor. We have

$$\chi_{ii} = \beta \left(\langle s_i^2 \rangle - \langle s_i \rangle^2 \right) = \beta \left(1 - \langle s_i \rangle^2 \right), \quad (1.35)$$

which means that the “local susceptibility” is

$$\chi_{\text{loc}} = \frac{1}{n} \sum_i \chi_{ii} = \beta \left(1 - q_{\text{EA}} \right), \quad (1.36)$$

where the first term β is the ordinary Curie susceptibility. This is equivalent, up to a prefactor, to the experimentally measurable susceptibility $\chi \propto \sum_{ij} \chi_{ij}$, provided $P(J_{ij}) = P(-J_{ij})$ (see next paragraph) and $h = 0$. This also provides a way to experimentally measure the EA parameter.

In order to prove the existence of a SG phase one can try to estimate q_{EA} and show that it is nonzero in a certain region of parameters, while simultaneously $m = 0$.¹² This was first done by Edwards and Anderson with the Hamiltonian (1.30). Computing q_{EA} in this model for a generic choice of parameters $\{J_{ij}\}_{ij}$ is unfortunately impossible. A general strategy for dealing with disordered systems is to assume that in a large enough sample the disordered couplings $\{J_{ij}\}_{ij}$ can be modeled as i.i.d. random variables with a given distribution $P(J)$. Then, in the limit of very large system sizes we expect almost any

¹¹Conceptually, fixing a component may be accomplished by setting up a small nonuniform magnetic field $h = (h_0, \dots, h_{n-1})$ and letting $\|h\| \rightarrow 0$ at the end, in analogy with the FM case; in practice, however, unlike in ferromagnets one is unable to determine the right choice of magnetic field because that requires the *a priori* knowledge of all local magnetizations $\langle s_i \rangle$, so different approaches are used instead [184].

¹²By the Cauchy–Schwartz inequality, q_{EA} is trivially nonzero whenever $m \neq 0$.

choice of parameters to yield the same value for intensive observables — a feature known as *self-averaging*. Moreover, we assume that the values of the $\{J_{ij}\}_{ij}$ are fixed, *i.e.* they do not measurably change over the observational time scale. This is done in the spirit of the Born–Oppenheimer approximation and will be taken as valid in all our examples where disorder plays a role. We say that the random couplings are *quenched* (non-dynamical) as opposed to *annealed* (dynamical).

In the above assumptions, we can get rid of the detailed parameter dependence by performing a *disorder average* over the parameter choice, which we denote by double brackets:

$$\llbracket f \rrbracket = \int \left(\prod_{ij} P(J_{ij}) dJ_{ij} \right) f(\{J_{ij}\}_{ij}). \quad (1.37)$$

The computation of the disorder-averaged $\langle s_i \rangle$ or $\langle s_i \rangle^2$, as well as other self-averaging quantities, proceeds through the calculation of the free energy,¹³

$$\llbracket f \rrbracket = -\frac{1}{\beta n} \llbracket \log Z \rrbracket, \quad (1.38)$$

which in turn requires computing $\llbracket \log Z \rrbracket$. The calculation is made possible by the famous *replica trick* for the partition function:

$$\log Z = \lim_{m \rightarrow 0} \frac{Z^m - 1}{m}, \quad (1.39)$$

where m is treated as an integer throughout the calculation to make the computation of $\llbracket Z^m \rrbracket$ simpler — as it can now be regarded as the partition function of m noninteracting replicas of the original system. The number of replicas m is then sent to 0 at the end with a process of analytical continuation.

In the case of the EA model, the replica calculation can only be seen through in the mean-field approximation, leading to results which have some qualitative agreement with the experiments (and some disagreements as well). Similarly, the free energy can be computed exactly if the model is genuinely mean-field, namely if the sum in Eq. (1.30) is replaced by a sum over all spin pairs and the couplings J_{ij} are appropriately rescaled to keep the energy extensive (*Sherrington–Kirkpatrick model* [229]). In both cases, the existence of a SG phase can be established, in the sense that $m = 0, q_{\text{EA}} \neq 0$, at low temperature. However, some unphysical results such as a negative value of the entropy are present in the theory as well.

¹³For instance, by adding a $h_i s_i$ term to the Hamiltonian and computing the corresponding free energy (which can be done in the EA model), one simply has $\llbracket \langle s_i \rangle \rrbracket = n \left. \frac{d \llbracket f \rrbracket}{d h_i} \right|_{h=0}$, where the disorder average passes through the derivative as h_i is disorder-independent.

The problem of negative entropies was dealt with by Thouless, Anderson and Palmer by an application of the so-called ‘Bethe method’ (today known as *cavity method*), leading to the celebrated TAP equations [251]. Essentially, they improved the mean-field equation for $\langle s_i \rangle$ by an extra local-field term which is typically discarded during the approximation process, being only $O(1/n)$ in ferromagnetic models (but is in fact $O(1)$ for spin glasses).

Along with the formulation of the TAP equations came the realization of the importance of “rugged energy landscapes” in determining the low-temperature properties of glassy systems. Indeed, it was soon discovered that the failure of the mean-field calculation to yield sensible results in the deep spin-glass phase of the SK model was due to the wrong assumption that all “replicas” of the system introduced by formula (1.39) would behave identically. More precisely, in the determination of the free energy of the SK model, one introduces dummy integration variables $\{q_{\alpha\beta}\}_{\alpha<\beta}$ in order to get rid of the biquadratic part through the well-known Gaussian trick

$$e^{k^2/2} = \int_{-\infty}^{+\infty} \frac{dx}{\sqrt{2\pi}} e^{-x^2/2+kx}, \quad (1.40)$$

the α, β indices run over the m replicas, as there is one biquadratic term per pair of replicas. Upon substitution (1.40), one is left with a formula which is amenable to steepest descent in the n variable, which fixes the values of the dummy variables to their optimum point. The free energy ends up depending on the sum $\sum_{\alpha<\beta} q_{\alpha\beta}^2$, where $q_{\alpha\beta}$ is now the optimum value. Sherrington and Kirkpatrick, arguing from the indistinguishability of the replicas, had assumed that $q_{\alpha\beta} = q$ for all $\alpha \neq \beta$ (*replica symmetry*). This solution, however, despite always being a stationary point for the integrand of the free energy, fails to be a maximum point (as required by the steepest descent method) once T is lowered past the so-called de Almeida–Thouless instability line [75]. In other words, the solution can be improved by a non-uniform choice of $\{q_{\alpha\beta}\}$ coefficients, signaling a *replica symmetry breaking* (RSB) in the system. The “ 0×0 ” symmetric matrix¹⁴ $(q_{\alpha\beta})$ (with $q_{\alpha\alpha} = 0$) must be interpreted as the true spin-glass order parameter: the replica-symmetric ansatz corresponds to setting $q_{\alpha\beta} = q$ for all $\alpha \neq \beta$. In that case, q coincides with the Edwards–Anderson parameter, giving us the alternative interpretation of q_{EA} as the average overlap between two distinct replicas. In the paramagnetic phase, where replica symmetry holds, two independent replicas have spins pointing in essentially uncorrelated directions, resulting in a subextensive overlap ($q_{EA} = 0$). At lower temperatures, pure states start to play a role and RSB must be accounted for.

An explicit formulation of the RSB solution for the SK model was worked out by Parisi [185,200,201], who showed how to properly parametrize the $(q_{\alpha\beta})$ matrix to account for the breaking of the index-permutation symmetry. This is done in a hierarchical fashion:

¹⁴The terminology is to be understood as the $m \rightarrow 0$ analytic continuation of a parametrized space of $m \times m$ matrices.

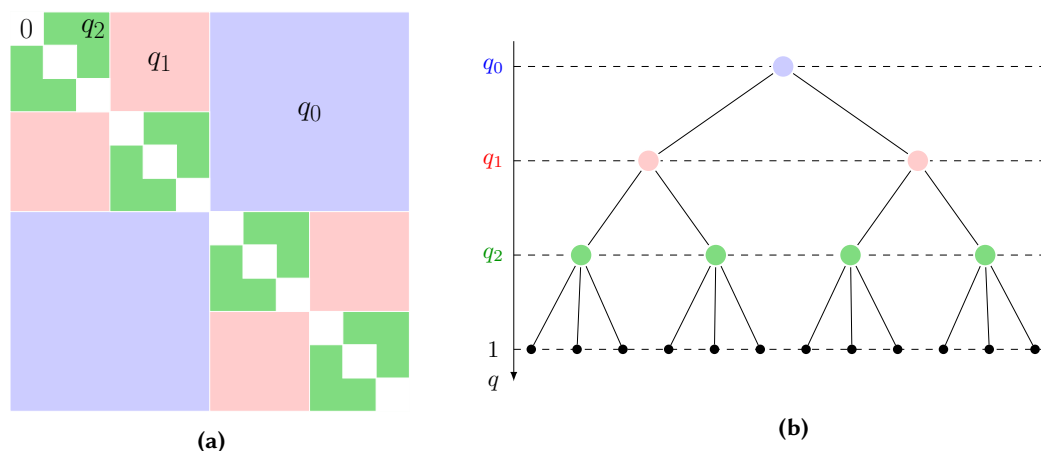


Figure 1.2: (a) The Parisi spin-glass order parameter is an $m \times m$ symmetric matrix $(q_{\alpha\beta})$ (analytically continued to $m \rightarrow 0$) with a hierarchical structure determined by parameters k_1, \dots, k_r (see main text). Here is an example with $m = 12, k_1 = 6, k_2 = 3$. Full RSB requires an infinite number of parameters. (b) The hierarchical structure of $(q_{\alpha\beta})$ is the manifestation of an underlying clustering of states, with the overlap between states only depending on their graph distance.

at the simplest level, introduce a parameter k_1 and use the following ansatz:

$$q_{\alpha\beta} = \begin{cases} q_1 & \text{if } \lfloor \alpha/k_1 \rfloor = \lfloor \beta/k_1 \rfloor, \\ q_0 & \text{otherwise.} \end{cases} \quad (1.41)$$

The free energy must then be optimized with respect to k_1 , subject to $k_1, m/k_1 \in \mathbb{N}$. This is called the *1-step replica symmetry breaking* (1RSB) ansatz. More generally, one can introduce several parameters $k_1 \geq \dots \geq k_r$ and set $q_{\alpha\beta} = q_i$ if $\lfloor \alpha/k_j \rfloor = \lfloor \beta/k_j \rfloor$ is satisfied for $j \in \{1, \dots, i\}$. All the k_j should be integers, as well as their consecutive ratios — this corresponds to a partitioning of $[m]$ into progressively larger “containers” (see Fig. 1.2a).

The r -step RSB ansatz is therefore characterized by a collection of r commensurate integers, $\{k_i\}_{i=1}^r$, and $r + 1$ values $\{q_i\}_{i=0}^r$ for the replica–replica overlaps, which together can be encoded in a function

$$q(x) = q_i \quad \text{if } k_{i+1} \leq x < k_i, \quad (1.42)$$

with $1 \leq x \leq m$. In the $m \rightarrow 0$ limit, one has $0 \leq x \leq 1$ instead, with a corresponding inversion in the order relation of the k_i 's. This can be generalized to the *full replica symmetry breaking* ansatz, $r \rightarrow \infty$, by trading Eq. (1.42) for a continuous function $q : [0, 1] \rightarrow [q_{\min}, q_{\max}] \subset [0, 1]$ functioning as an infinite-dimensional order parameter for the SG transition.

The breaking of replica symmetry is the mathematical manifestation of a rugged energy landscape. Indeed, one can think of the Parisi ansatz for $(q_{\alpha\beta})$ as reproducing an underlying structure where configurations of similar energies organize themselves into clusters.¹⁵ In the low- T solution, each replica independently probes different configurations, which need not belong to the same cluster. Moreover, clusters themselves may contain subclusters in a hierarchical fashion, so that the overlap between two configurations depends on how many layers of clustering one has to traverse in order to connect them. This layered structure is conveniently represented as a tree (Fig. 1.2b), where the distance between two configurations (defined as $d(\alpha, \beta) = 1 - q_{\alpha\beta}$) only depends on the level of their nearest common ancestor [85].

This high-level understanding of spin-glass energy landscapes in terms of clusters of states has been extremely prolific. For instance, it underlies the idea that low-energy states are to be found at the bottom of a large number of “valleys” separated from each other by tall free-energy barriers, which in turn motivates and explains the success of powerful algorithms such as simulated annealing and quantum annealing (which will be described in the next section). As already mentioned, the most impressive feature of spin glasses lies in their horizontal reach, with similar ideas echoing through all sorts of different disciplines where optimization is at work. Computational techniques that are useful for estimating low-temperature properties of spin glasses can also be applied to these disciplines, and vice-versa.

In the next section, which wraps up our introduction, we briefly discuss one of the most important stochastic algorithms that have been devised for this very purpose — namely, tackling optimization problems — in the novel framework of quantum computation. One of the flagship goals of quantum computers is their ability to efficiently simulate complex physical systems. This includes the potential of studying glassy problems, and solving (or at least approximating) hard optimization problems of much larger sizes than what is classically achievable. The comparison between classical and quantum algorithms will be the topic of Chapters 3 and 4.

1.3 Computing with quantum mechanics

1.3.1 A foreword on quantum computation

The birth of quantum computation (QC) is traditionally¹⁶ dated back to Feynman’s remark [103] on the inadequacy of using classical physics, which most of the technology empowering traditional computers is based on, in order to simulate quantum systems. The implications of a quantum computational machine were first discussed by Albert [13]

¹⁵Similar clustered phases are known to exist in optimization problems of the *constraint satisfaction* type (CSPs), such as k -SAT and q -COL [33].

¹⁶Soviet physicist Manin had conceived of similar ideas in his earlier book [168] (cit. in Ref. [231]).

and then more thoroughly by Deutsch [82], who formulated a theory of quantum Turing machines.

The true potential of QC emerged with the discovery of algorithms that could outperform their classical counterparts either provably, such as Deutsch and Josza's [83] or Grover's [116] algorithms, or practically, such as Shor's algorithm [230] for factoring integers.¹⁷

Presently, two routes have proved promising toward the implementation of useful QC, already leading to the creation of extremely sophisticated (if rather impractical) machines. These are the *circuit model* of QC and the *quantum annealing* paradigm, which will be discussed in Section 1.3.2. The circuit (or gate) model may be considered the standard model of quantum computation, for both historical and logical reasons. It is the most natural translation of traditional electronic circuits into a quantum setting, where bits are replaced by two-level quantum systems (*qubits*) and logic gates are implemented via unitary operators. It is relatively straightforward to demonstrate that circuit-based QC is *universal*, in the sense that any unitary transformation on any number of qubits can be realized with arbitrary accuracy, given a small toolbox of simple gates [194]. The physical realization of a quantum circuit can be accomplished in a number of different ways, including optical systems, ion traps, and solid state technology. The current frontier of research appears to be superconducting qubits, which can be arrayed on an integrated circuit and manipulated via tunable electromagnetic fields. The degree of quantum control achieved in systems of over 50 qubits of this type is so high already that results were published claiming quantum supremacy in specific tasks [26], though the claims were later challenged.

Although quantum supremacy — meaning the complete obsolescence of classical computers in performing certain useful tasks — is far from established, quantum computers and quantum technologies in general (including quantum communication, quantum sensors, quantum cryptography and more) are expected to evolve quite rapidly in the near future and eventually enter our daily life in a revolution comparable to the advent of the internet at the turn of the last century (see *e.g.* [79]). Even today, so-called *noisy intermediate-scale quantum* (NISQ) devices are already starting to show the promise of near-term applications, if not for universal computation, at least as special-purpose machines (in particular for the simulation of modest-size quantum systems) [211].

Obviously the topic of quantum computation, whether the theory or the history or even just recent developments, is far too wide to be properly contained in this introduction. The interested reader will find in Ref. [194] a classic textbook, and in Ref. [211] a good review of the current state and expectations for NISQ computation. In the rest of the section,

¹⁷From a complexity-theory standpoint, the class BQP of problems efficiently solved by quantum computers is known to be intermediate between P (problems solvable in polynomial time) and PSPACE (problems solvable with polynomial memory). Unfortunately, it remains to be proven that PSPACE is actually larger than P. Even assuming that $P \neq NP$ (problems admitting a poly-time verifiable certificate), the relation between NP and BQP remains unclear. It is possible that neither class fully contains the other [194].

however, we find it useful to present and discuss the concept of quantum annealing, as this will ensure a more thorough understanding of Chapter 3.

1.3.2 Quantum annealing

Annealing is the process of heating up and then cooling down a sample (typically a metal) in order to affect its structural properties in some specific way. The idea of cooling a system slowly enough to let it accommodate into a lower-energy state than it would otherwise, already used in practice *e.g.* to grow crystalline structures from a melted material, was appropriated by theoreticians and applied to the context of computer simulations in the 1980s. The term *simulated annealing* (SA), in particular, is found in a paper by Kirkpatrick and others [144], where they successfully applied it to the classic traveling salesman problem. It is worth mentioning that Kirkpatrick stumbled upon these ideas as a consequence of his work on spin glasses, *cf.* Section 1.2.

In SA, one is interested in minimizing some Hamiltonian H encoding a complex optimization problem. Although the idea at the basis of SA is general enough to be viewed as a metaheuristic, one of its most basic formulations is in terms of a typical Markov-chain Monte Carlo simulation: one initializes the system in a random configuration and runs a stochastic (*e.g.* Glauber or Metropolis) dynamics at some “high” initial temperature T_0 , which will cause the system to sample states according to the Gibbs distribution

$$P_{T_0}(x) = \frac{e^{-H(x)/T_0}}{Z_{T_0}}. \quad (1.43)$$

Then, the temperature is gradually decreased down to some final value $T_1 \ll T_0$, reducing thermal fluctuations and concentrating the Gibbs measure on low-energy configurations. Ideally, one would like the system to eventually reproduce the new thermal distribution $P_{T_1}(x)$, but in practice this can almost never be achieved. The reason should be clear by now: since optimization problems possess rugged energy landscapes, their thermalization requires prohibitively long times, meaning that the effective distribution observed from the simulation will at best approximate a restricted ensemble during the available time frame (see Section 1.1.3).

Therefore, lowering the temperature will cause the simulation to eventually “freeze” into some pure state, similar to what would happen by simply starting the simulation from the “cold” temperature T_1 . The major difference is that while in the latter case the system would immediately get stuck upon falling into the local minimum point closest to the initial state, in SA the system is given the chance to first explore (at least in principle) the entirety of the configuration space during its “hot” phase, and as its thermal energy is gradually depleted, it will be more likely to settle into a configuration more closely matching the global energy minimum, as relatively shallow local minima can be escaped at earlier stages. This process of cascading of the distribution requires the annealing schedule to be performed slowly enough, which introduces a tradeoff between the quality of the final

solution and the time required by the simulation, as should be expected of any optimization algorithm. With its combination of effectiveness and conceptual simplicity, SA remains nowadays one of the most popular classical algorithms for general-purpose optimization, and its applications are widespread both in the sciences and the industry [225].

Quantum annealing (QA) may be viewed as a variant of SA where quantum fluctuations replace thermal ones as the driving mechanism for the dynamics. The basis of QA is the *adiabatic theorem* [16, 53], a standard result in quantum mechanics stating that, roughly speaking, a system prepared in the ground state of a Hamiltonian H_0 which is then changed in a continuous, time-dependent fashion, $H = H(t)$, will naturally time-evolve into the instantaneous ground state of $H(t)$ throughout the evolution so long as this occurs slowly enough.

In practice, the theorem admits several, inequivalent variants, so one should rather talk about *adiabatic theorems* (cf. Ref. [11]). These variants can be more or less rigorous, and assume more or fewer conditions on $H(t)$, with ensuing adiabatic bounds of different strengths. The most familiar and most quoted “adiabatic condition” [179] can be stated as follows: consider a time-dependent Hamiltonian $H(t)$, with the time variable ranging in $t \in [0, \mathcal{T}]$, and assume that its instantaneous spectrum $\{E_a(t)\}_a$ is nondegenerate for all t . This implies in particular that the order relation between the eigenvalues is preserved throughout the evolution,

$$E_0(t) < E_1(t) < \dots \quad \forall t \in [0, \mathcal{T}]. \quad (1.44)$$

Because of the nondegeneracy, we may unambiguously label $|E_a(t)\rangle$ the unique instantaneous eigenstate with energy $E_a(t)$ at time t . We now focus on the ground state $|E_0(t)\rangle$, though the theorem could be applied to excited states as well. Define the instantaneous ground gap

$$\Delta(t) = E_1(t) - E_0(t), \quad (1.45)$$

as well as the transition rate

$$R(t) = |\langle E_0(t) | \dot{H}(t) | E_1(t) \rangle|, \quad (1.46)$$

where $\dot{H} = \frac{dH}{dt}$.

The adiabatic theorem states that with high probability, the state resulting from time-evolving $|E_0(t=0)\rangle$ up to $t = \mathcal{T}$ will be found to be $|E_0(\mathcal{T})\rangle$ provided

$$\max_{0 \leq t \leq \mathcal{T}} \left[\frac{R(t)}{\Delta^2(t)} \right] \ll 1. \quad (1.47)$$

The qualifier “with high probability” means that the probability of the event *not* happening is upper bounded by the left-hand side of Eq. (1.47). An alternative formulation (which also avoids some criticisms of the Messiah criterion (1.47) not working for Hamiltonians

with externally-driven oscillations) was provided by Amin [18], who stated the following adiabatic condition:

$$\mathcal{T} \gg \max_{0 \leq s \leq 1} \left[\frac{R(s)}{\Delta^2(s)} \right], \quad (1.48)$$

where all the quantities are now written in terms of the normalized time $s = t/\mathcal{T}$ and the Hamiltonian is required to depend on t and \mathcal{T} only through s (*i.e.* there is a single time scale). The quantity $\epsilon = \mathcal{T}^{-1} \max_s [R(s)/\Delta^2(s)]$ upper bounds the difference $|1 - \langle E_0(\mathcal{T}) | \psi(\mathcal{T}) \rangle|$ (with $|\psi(\mathcal{T})\rangle$ the final state), so the probability of observing the final state in the ground state of the final Hamiltonian is at least $1 - \epsilon^2$.

It should be noted that neither condition (1.47) nor (1.48) are rigorous mathematical statements. Rather, they are necessary relations that should hold between \mathcal{T} and $H(t)$ in order for the evolution to be considered adiabatic. Nonetheless, the *inverse-square dependence* on the gap of the adiabatic time has essentially earned the status of “folk theorem” among physicists. More rigorous criteria exist providing rigorous upper bounds on *e.g.* the distance between the final state and the ground state of $H(\mathcal{T})$ (in the sense of operator norm of the difference between the respective density matrices), which tend to provide looser bounds. For instance, if $H(s) \in C^2([0, 1])$ and $\Delta(s) > 0$, a *sufficient* adiabatic condition is the following [11]:

$$\mathcal{T} \gg \max_{0 \leq s \leq 1} \left[\frac{\|\ddot{H}(s)\|}{\Delta^2(s)} \right], \max_{0 \leq s \leq 1} \left[\frac{\|\dot{H}(s)\|^2}{\Delta^3(s)} \right], \max_{0 \leq s \leq 1} \left[\frac{\|\dot{H}(s)\|}{\Delta^2(s)} \right], \quad (1.49)$$

where obviously $\ddot{H} = \frac{d^2H}{dt^2}$. Notice the inverse *cubic* dependence on the ground gap.

Another theorem provides an inverse-square bound (corrected by a $\Theta(|\log \Delta|^{6\alpha})$ term with $\alpha > 1$) for sufficiently well-behaved Hamiltonians. This bound is close to optimal as there is a *lower* bound on \mathcal{T} which has the same scaling save for the logarithmic correction, which is $\Theta(|\log \Delta|^{-1})$ [91]. We will not discuss the precise form of the adiabatic bound any further. In Chapter 3, we will adopt the traditional assumption that adiabatic evolution requires a time $\mathcal{T} \sim 1/\Delta^2$ to succeed with good probability. Logarithmic or even polynomial corrections will not alter the gist of the discussion, as will be clearer shortly.

Now that we have established the content of the adiabatic theorem, we can formulate the strategy underlying quantum annealing. Suppose we are dealing with an optimization problem, and that we are able to encode this problem as the Hamiltonian of some system. We call it H_p , the “problem Hamiltonian”. It is worth stressing that every ground state of H_p corresponds to a solution of the encoded problem. Typically, when we are dealing with a discrete configuration space (*combinatorial optimization*), H_p will be cast as a spin Hamiltonian.¹⁸ For instance, let $E : \{0, 1\}^n \rightarrow \mathbb{R}_0^+$ be the cost function of a combinatorial

¹⁸The hardness of the encoding will strongly depend on the locality structure of the problem, as the most general optimization problems will result in completely nonlocal Hamiltonians, which are hardly possible to construct. For now, we can suppose that we are somehow able to build H_p in our laboratory.

optimization problem defined on n -bit strings. The problem Hamiltonian will be of the form

$$H_p = \sum_{s \in \{0,1\}^n} E(s) |s\rangle \langle s|, \quad (1.50)$$

where $\{|s\rangle\}_{s \in \{0,1\}^n}$ is some orthonormal basis of the Hilbert space. Being the basis on which the problem Hamiltonian is diagonal, it is called the *computational basis*, and it represents the *classical* configuration space (for this reason, its elements ought to be product states). Typically, one identifies bits 0 and 1 with spin- $\frac{1}{2}$ states of definite z -component, $|\uparrow\rangle$ and $|\downarrow\rangle$, and then the computational basis is given by the eigenstates of the z -component of the total spin operator.

We can now initialize our spin system in the ground state of *another* spin Hamiltonian H_d . This should satisfy two crucial requirements: first, its ground state should be known and easily prepared; second, it should not be an eigenstate of the problem Hamiltonian as well (in particular, H_d and H_p should not commute). A typical choice of “driver Hamiltonian” is a transverse field term,

$$H_d = -\Gamma \sum_{i=0}^{L-1} s_i^x; \quad (1.51)$$

in this way, we have $[H_p, H_d] \neq 0$ and the initial state is simply the uniform superposition of all basis states. For the second part of the algorithm, we perform an adiabatic evolution of the initial state by slowly changing the Hamiltonian from H_d to H_p :

$$H(s) = A(s)H_p + B(s)H_d, \quad (1.52)$$

where $A(0) = B(1) = 0$ and $A(1) = B(0) = 1$. The pair of continuous functions $(A(s), B(s))$ is the annealing schedule of the algorithm, and is arbitrary as long as the aforementioned initial and final conditions are satisfied.

From the adiabatic theorem we know that, if the evolution is slow enough, the system will have a large superposition with the ground state of the problem Hamiltonian H_p , meaning that a measuring protocol in the computational basis will provide with high probability a solution to the encoded optimization problem. The process just described is called the *quantum adiabatic algorithm*, first described in Ref. [100].

In realistic setups, one does not expect the system to start from an *exact* ground state, much less to remain in one, due to the inevitable presence of thermal fluctuations and parametric errors (including errors in the transverse field strength and in the implementation of the annealing schedule). The nonzero-temperature version of the quantum adiabatic algorithm is what we refer to as quantum annealing algorithm (although the terms are often used interchangeably).

We can now see the analogies between QA and simulated annealing, described at the beginning of this section. Both QA and SA are optimization algorithms (or heuristics)

which exploit physical principles to explore the classical energy landscape. A typical, if overly simplified, picture of the situation is to think of the configuration space as a rugged environment of deep valleys or basins separated by tall cliffs or mountains. In SA, we have in mind a particle pushed across this environment by a stochastic dynamics, while in QA we have a quantum particle subject to the Schrödinger equation; in both cases, the “landscape” itself should be regarded as some effective (semiclassical, in the case of QA) potential, and the particle is trying to find a (hopefully global) minimum point where to settle. To this end, SA relies on thermal fluctuations to help the particle climb the barriers, whereas QA uses quantum fluctuations to tunnel through them.

Taking this picture seriously, perhaps too much so, one can estimate the rate at which the stochastic and quantum walkers are able to overcome an energy barrier ΔE (or more accurately, a free energy barrier ΔF). In the case of SA, the characteristic time of thermal activation takes the Arrhenius form

$$\tau_{\text{SA}} \sim e^{\Delta E/T}, \quad (1.53)$$

while for QA one can use the WKB estimate

$$\tau_{\text{QA}} \sim e^{w\sqrt{m\Delta E}/2} \quad (1.54)$$

for a barrier of width w [158]. As a result, one expects QA to outperform SA whenever the energy landscape is characterized by “tall, thin barriers”, *i.e.* when the combination $w\sqrt{\Delta E}$ is typically less than ΔE (assuming the remaining parameters can be neglected in the economy of the argument). This kind of reasoning should be taken with a grain of salt, offering at best a qualitative suggestion in support of QA. For instance, the Arrhenius and WKB estimates for the characteristic time scales of the algorithms are strictly valid only in the one-dimensional case, which could not be further removed from the n -dimensional landscape of the actual optimization problem. The problem of directly comparing SA with QA is a complicated one and remains outstanding, nor do we have a clear-cut and general understanding of which classes of problems we should expect to be better suited for one as opposed to the other.

To date, Canadian company D-Wave Systems is the only manufacturer to have produced and marketed a quantum annealer (2048 superconducting qubits as of the D-Wave 2000Q, the latest public version), and although many studies have been performed on it looking for quantum advantage (see *e.g.* [80]), this has so far proved an elusive quest [166, 167]. As a general principle, a number of limitations of adiabatic quantum computing (and *a fortiori* QA) must always be kept in mind. First off, unlike the circuit model of QC, quantum annealers are not currently conceived as general-purpose machines, but rather as optimizers. Although there are theorems proving the universality of adiabatic QC (indeed, it is equivalent to circuit-based QC up to a polynomial overhead [9]), the simulation of arbitrary quantum circuits by a quantum annealer is out of the scope of a device like the D-Wave 2000Q. Secondly, the quantum control on annealers is not

nearly on par with circuit-based quantum processors, mostly due to the relatively huge number of qubits (over 2000¹⁹ versus less than 100). In general, there are serious doubts about the scalability of quantum annealers with respect to qubit fidelity, as the theory of error correction for this kind of systems is still non-existent. Still, the development of larger and (more importantly) less noisy annealers in the near future [176] is expected to improve our understanding on these matters, and possibly usher in an era of quantum advantage for such problems as the simulation of Hamiltonians with no “sign problem” (see Section 3.2.1) [211].

¹⁹It should be noted that this is the number of physical units in the annealer, which is not the same as the number of *logical* qubits that can be processed. A logical qubit requires several (say, 5–10) physical qubits to represent, shrinking the representable configuration space. A similar unit reduction happens in circuit-based models when applying error-correcting codes.

Chapter 2

Many-Body Localization and Symmetry

Ἄρμονιᾳ ἀφανῆς φανεροῦς κρείττων.

A concealed harmony is superior to a visible one.

Heraclitus of Ephesus
cit. in St. Hippolytus of Rome
Refutatio omnium heraesium IX, 9 [169]

2.1 Localized quantum systems

In Section 1.1 we discussed the fact that equilibrium statistical mechanics is ultimately founded on the ergodic hypothesis, which is most aptly expressed in the quantum theory by the eigenstate thermalization hypothesis (ETH), Eq. (1.18). We also saw how ergodicity can be (trivially or nontrivially) violated in both classical and quantum systems in the presence of conserved quantities or structural disorder. Classical spin glasses in particular were discussed in some depth (Section 1.2), as they constitute the first and best understood example of highly non-ergodic systems in statistical physics.

In this chapter, we turn our attention toward a different, purely quantum mechanism of ergodicity breaking, *many-body localization* (MBL). Similar to glassy systems, it involves the *localization* in space of some degree of freedom, but unlike structural and spin glasses, MBL is not caused by kinematic constraints from “rugged” free energy landscapes, but rather by the effect of quantum interference. More explicitly, recall that glassiness can be construed as a form of extreme slowdown in a system which is trying to thermalize, but prevented from doing as a consequence of vanishingly small escape rates from some component of the configuration (or Hilbert) space, in which the system finds itself trapped. Given enough time, a stochastic (thermal or quantum) fluctuation would eventually allow

any such (finite) system to attain its true Gibbs distribution, “melting” the glass.

In contrast, as we will see shortly, MBL should be viewed as a form of *integrability* in the quantum sense. Namely, in an MBL system one can identify an extensive ($O(n)$ for an n -particle system) set of exactly conserved charges, which are crucially *quasilocal* operators, *i.e.*, whose support is exponentially bounded away from a particular point in space. Similar to the case of ordinary integrability, the proliferation of conservation laws constrains the dynamics to such a degree that the stationary state of such a system is no longer thermal. This remains true even in the thermodynamic limit, where it can be proven that the usual picture of “large system acting as a bath for itself” must fail in the MBL phase. The system instead retains *local* information about a generic initial state for an infinitely long time.

Another important feature setting the MBL regime apart from glasses is that the localization transition is not accompanied by any thermodynamic signature, such as the appearance of a nonzero Edwards–Anderson parameter, as it would be the case for a spin glass. Rather, it represents a *dynamical phase transition* whose presence can only be assessed by looking at spectral or nonequilibrium properties of the system. In particular, the ground state physics is hardly affected by the onset of MBL, which, unlike ordinary quantum phase transitions, is manifested at high energy densities (or equivalently, high temperatures).

In the following section, we give an introductory glance at the single-particle version of MBL, first described by P. W. Anderson in his authoritative 1958 paper [19] and appropriately called *Anderson localization*. In the same paper, Anderson suggested that this phenomenon should survive the introduction of sufficiently weak interactions between particles, though the problem remained open. Physicists had to wait almost half a century for a full answer, which came in the form of a comparably foundational article by Basko, Aleiner and Altshuler in 2006 [35]. This will be the matter of Section 2.1.2.

2.1.1 Anderson localization

The Anderson model is the ensemble of tight-binding Hamiltonians

$$H = -t \sum_{\langle ij \rangle} (c_i^\dagger c_j + c_j^\dagger c_i) + \sum_i \epsilon_i c_i^\dagger c_i, \quad (2.1)$$

where the $\{c_i\}_i$ operators are annihilation operators of fermionic character and $\{\epsilon_i\}_i$ are quenched i.i.d. random variables, whose distribution is typically taken to be the uniform one on some symmetric interval $[-W/2, W/2]$. We may always set $t = 1$ by an appropriate choice of energy units. W then parametrizes the *disorder strength*. The Anderson model can be defined on any graph, and the first sum is taken over its edge set. We will talk about the Anderson model “in d dimensions” by taking the underlying graph to be the square grid lattice with coordination number $2d$ (linear chain, square lattice, *etc.*).




Figure 2.1: The Anderson model: fermions can hop to unoccupied neighboring sites with amplitude t , while experiencing a random on-site potential varying on a scale W . The ratio W/t (disorder strength) controls the dynamical regime of the model: a small W/t results in diffusive behavior for the particles, while for large disorder the particles are exponentially localized. In $d \leq 2$, the model is always localized.

Hamiltonian (2.1) is a model of free particles hopping on a lattice with on-site disorder (Fig. 2.1). In the clean case, $W = 0$, it is known that the eigenstates of the system are simply plane waves, and in particular, they are fully delocalized on the lattice. The same would be true if the on-site potential were replaced by a periodic potential, with the creation of Bloch states.

It was therefore surprising to learn that, in the one-dimensional case, the Anderson Hamiltonian admits *no* extended¹ eigenstates at *any* nonzero value of the disorder. Indeed, it can be proven that any disorder $W > 0$ makes the spectrum of H entirely pure-point and all the corresponding eigenfunctions exponentially bounded in space with probability one (*exponential spectral localization*). Not only that, but a stronger, *dynamical localization* condition can also be proven, ensuring that all spatial moments of any compactly supported initial condition $|\psi(t=0)\rangle$ (namely $\langle \psi(t) || x|^p | \psi(t) \rangle$ for all $p > 0$) stay finite at all times $t \geq 0$ [124, 243].

The situation is more interesting in higher dimensions. The two-dimensional Anderson model is something of a borderline case, corresponding to the critical dimension of the scaling theory [6]. It has fully localized spectrum in the strict sense, but delocalization can be induced by introducing perturbations such as a weak spin-orbit coupling [94]. We discuss instead the three-dimensional model, which presents a perfectly clear example of

¹To be clear, here and everywhere else (including Ch. 4) we use the terms “delocalized” and “extended” interchangeably to denote a quantum state whose support over the real space is “large”, *i.e.*, for which the number of basis states required to approximate it to within some given precision diverges in the infinite-size limit, when using the natural product basis determined by the lattice structure.

localization–delocalization transition as described hereunder. We should point out that, unlike in the 1D case, the following discussion mostly relies on heuristic evidence rather than mathematical rigor.

In $d = 3$, Anderson localization (AL) is only observed for sufficiently large values of the disorder, whereas in the low-disorder regime all the eigenfunctions are fully extended, analogously to the clean model. Moreover, the transition between fully extended and fully localized phases occurs in a gradual manner as W is increased: at intermediate values of the disorder, *some* eigenstates are localized while others are not. In particular, there is a certain interval of energies $I = I(W) \subset \text{sp}(H)$ in the spectrum such that eigenstates with associated energies $E_a \in I$ are delocalized, while the rest are localized. In the infinite-size limit, the spectrum is dense and bounded,² so the interval $I(W)$ can be uniquely defined. Its boundary ∂I , separating the two regions, is called the *mobility edge*.

The interval straddles the middle of the spectrum, so that the “easiest” eigenstates to delocalize are, perhaps unsurprisingly, the infinite-temperature ones. The disorder strength required for a nonempty mobility edge to appear is deterministic in the large-size limit and represents a critical disorder,

$$W_c = \min\{W \mid I(W) = \emptyset\}. \quad (2.2)$$

Its value, which is zero in 1D and 2D, has been estimated at $W_c \approx 16.5$ in 3D [186,233] and grows with the dimensionality [245].

The appearance of a localizing phase at $W = W_c$ is called the *Anderson transition* (AT) or the *metal–insulator transition* (MIT).³ The latter term reflects the fact that the presence of localized states in the spectrum has a profound dynamical impact on the system. Namely, it entails the impossibility for a particle to traverse the system, and if the particle is interpreted as a charge carrier, the DC conductivity (see Eq. (2.20) below) and other diffusion-related coefficients must drop to zero as a consequence, making the system an insulator.

Of course, a theory of free particles makes for an extremely lackluster description of real materials. Still, the phenomenon of AL and the MIT in three dimensions seemed too good of a lead to discount, and as a consequence of Anderson’s paper, a remarkable amount of effort has been devoted to gaining a more thorough understanding of localizing systems.

Theoretical progress, though significant, remained mostly confined to the single-particle case (for a comprehensive review, see Ref. [150]). It was not until the 2000s that the problem of interactions was finally incorporated. Since then, huge strides have been made by the freshly born research community in many-body localization, including some impressive experiments that have been since made possible by the quick progresses

²The spectrum of Hamiltonian (2.1) in any dimension d is almost surely contained in the interval $[-2d - W/2, 2d + W/2]$.

³One typically uses the term “Anderson transition” to underline the single-particle nature of the model.

in cold-atom and trapped ion technology. We give a short summary in the next section, and present a more systematic description of MBL in Section 2.2.

2.1.2 The effect of interactions

The exponential localization of all the single-particle states ceases to be enough proof of absence of quantum transport once interactions between the particles are turned on. Indeed, the completely noninteracting nature of the Anderson model can be considered a saving grace by the theoretician seeking to rescue standard thermodynamics, as the ergodicity breaking present in the model can be attributed to such an unrealistic assumption.

For instance, it is well-known that thermally-activated phononic processes can induce hopping between localization centers in insulators via a process called variable-range hopping [190]. Including these processes into an otherwise fully AL model would inevitably induce a small ($O(e^{-c/T^a})$) conductivity, eventually restoring ergodicity at any $T > 0$.⁴ However, interactions between electrons are not limited to phonon-mediated processes, so it is natural to ask what influence the various electronic exchanges have on localization.

This question was tackled Basko, Aleiner and Altshuler [35] in a pioneering work (see also Ref. [115]), where they proved that a short-range, weakly-interacting model of electrons which is fully AL in the noninteracting limit displays both an insulating and a metallic phase depending on the disorder's strength. Moreover, the insulating phase has *strictly zero* conductivity even at nonzero temperature, in contrast with ordinary insulators. The stability of the two phases was proved to withstand loop corrections from all orders of the Feynman diagram expansion. The fact that different phases exist in different parameter regimes proves in turn that the model must undergo a phase transition (or a crossover) in between.

These works solidly established the existence of a *many-body localized* (MBL) phase in isolated disordered systems with weak interactions, and spurred new interest in the physics community to look into the properties of such a phase. Oganesyan and Huse [195] followed up with a numerical analysis of a disordered lattice model which aimed at showing how finite-size spectral statistics could be used as a witness of MBL (see Section 2.2.1). Of particular import in their work was an ingenious single-number diagnostics they devised in the form of the “*r*-parameter”, namely the average ratio of two consecutive level spacings, which was destined to become very popular in subsequent numerical work. However, the noticeable finite-size effects due to small system sizes ($L \leq 16$) prevented them from making conclusive statements. Indeed, a similar study of level spacing statistics had already been performed in a 2D system, and concluded that localized statistics could only hold for the ground state, *i.e.*, interaction would always delocalize the system at any $T > 0$ [236]. It is nowadays believed that this need not be the case.

After Oganesyan and Huse's work, many other litmus tests were designed to numerically assess the presence of a MBL phase in small systems — typically spin chains with

⁴Note that the DC conductivity is instead nonzero even at $T = 0$ for delocalized models.

some form of quenched disorder. Indeed, despite the limitations of exact diagonalization algorithms (often limited to $L \leq 20$ or even less for generic spin systems), a wealth of information has nevertheless been unlocked by such methods, as well as different numerical techniques such as the integration of the Lindblad equation (see *e.g.* [270] for a recent example). As we will see shortly, the time evolution of MBL systems is characterized by a slow propagation of entanglement, which makes it possible to apply certain numerical strategies, such as DMRG [264], taking advantage of matrix product state (MPS) representations of the quantum state [223] which are not generally available at nonzero temperatures.

The resulting phenomenology obtained from this kind of studies, as well as an explanatory framework in terms of emergent integrability, are going to be the subject of the next section.

2.2 MBL detectors and LIOMs

In this section, an overview is proposed of the various ways used to detect and characterize the many-body localized phase. Next, the unifying theory of local integrals of motion (LIOMs) is briefly presented. This constitutes the current reference frame for understanding the physical origins of the localization phenomenon, and it will be taken as a sort of “formal definition” of MBL. As we will see in Sections 2.3 and 2.4, this is not the only possible mechanism of interference-induced localization: different forms of nonthermal behavior are indeed observed in certain systems, which elude the LIOM description (at least to some degree). It is still an open and exciting question whether these forms of weak non-ergodicity can survive in the thermodynamic limit — thereby providing examples of novel phases of matter — or if they constitute transient (albeit long-lived) regimes which eventually give in to ergodicity.

2.2.1 Phenomenological characterizations of MBL

As is always the case for novel areas of study, research on MBL has first been aimed at collecting and categorizing the disparate evidence and phenomenology exhibited by putative MBL systems, whether by experimental or (for the most part) numerical investigation of a number of candidates. This resulted in a large quantity of characterizations and “detectors” of localization, all consistent with the same underlying idea of a lack of thermalization due to some hindrance in the transport dynamics of the system, which manifested itself both directly (vanishing conductivities, anomalous correlators) and indirectly (violation of ETH in eigenstates and nonthermal spectral statistics).

In this section, we explore some of the most important witnesses of MBL, without delving too deeply into the specifics, as a way of quickly savoring the topic. The interested reader is urged to look at more comprehensive reviews such as Refs. [4, 14, 192].

Eigenvalue and eigenstate statistics

Random matrix theory (RMT) was first introduced to physics by Wigner in the 1950s to explain the empirical distribution of the spacings between nuclear resonances [87, 265]. The idea of using *randomness* to reproduce deterministic, but *chaotic*, features of a theory proved very fruitful, and soon found its way into many other areas of physics and beyond.⁵ While a full review of the history and successes of RMT is far out of scope, we briefly recall the most relevant results for our current purposes. Readers looking for a deeper understanding will find in Mehta’s book [177] the golden-standard reference text.

Ergodic many-body systems are known to be very well approximated by random matrix ensembles, in the sense that their eigenvalues as well as eigenstates have the same statistical properties of those of a Gaussian random matrix which is only constrained by the symmetries of the system.⁶ This can be intuitively understood as a form of central limit theorem, where large matrices whose entries are essentially identically distributed (and whose moments are not fast-growing) eventually come to resemble Gaussian matrices in their spectral features. Hamiltonians of ordinary many-body systems tend to be of this type.

In particular, the random matrix ensemble that accurately describes time-reversal-invariant systems is the Gaussian orthogonal ensemble (GOE), namely, the ensemble of real symmetric⁷ $D \times D$ matrices $M = (M_{ab})_{ab}$ characterized by the joint pdf for their independent entries $\{M_{ab}\}_{a \geq b}$

$$P(M) \propto e^{-\frac{1}{4}\text{Tr}(M^2)}. \quad (2.3)$$

It is not hard to show that this is equivalent to a matrix whose independent elements are Gaussian distributed, with zero mean and variance 2 on the main diagonal, and zero mean and unit variance above it.⁸

In the large- D limit, the eigenvalue distribution (*i.e.* density of states) ρ for a GOE matrix tends with probability one to Wigner’s semicircle law,

$$\rho(\lambda) \sim \frac{\sqrt{4D - \lambda^2}}{2\pi D}. \quad (2.4)$$

Also, the distribution of *level spacings*, namely $\Delta_a = \lambda_{a+1} - \lambda_a$ (where $\lambda_0 \leq \dots \leq \lambda_{D-1}$ are the eigenvalues of the matrix), approximately tends to the *Wigner surmise*, more commonly

⁵Not unlike statistical physics, RMT takes full advantage of the phenomenon of “self-averaging”, which underpins the almost paradoxical fact that a system with an overwhelmingly large number of degrees of freedom is oftentimes easier to describe theoretically than a few-body problem.

⁶In particular, the Bohigas–Giannoni–Schmit conjecture [47] says that fluctuation laws of time-reversal-invariant quantum systems admitting a chaotic semiclassical limit are well described by the GOE ensemble.

⁷The denomination “orthogonal” does not refer to the matrices themselves, but rather to the invariance of the joint pdf (2.3) under orthogonal transformations. In other words, if a symmetric matrix M is GOE-distributed, so is $O^T M O$ for any orthogonal matrix O .

⁸In practice, a GOE matrix M is conveniently obtained by generating D^2 i.i.d. variables with zero mean and unit variance, arranging them into a matrix A , and then defining $M = \frac{1}{\sqrt{2}} (A + A^T)$.

referred to as *Wigner–Dyson distribution* in the MBL community:

$$P_{\text{WD}}(s) \sim \frac{\pi s}{2} e^{-\frac{\pi s^2}{4}}, \quad (2.5)$$

with $s = \Delta / \langle \Delta \rangle$ ($\langle \Delta \rangle = D^{-1} \sum_{a=0}^{D-1} \Delta_a$).

The property $P_{\text{WD}}(s \rightarrow 0) \rightarrow 0$ is called *level repulsion*, because it prevents two consecutive eigenvalues (energy levels, in the context of physics) from coming too close together. In particular, a random matrix has zero probability of having degenerate energy levels.

This tells us that, perhaps unintuitively, the spectrum of a random matrix does *not* resemble a collection of independent random variables, but rather, important correlations exist between the eigenvalues. Indeed, the level spacings of a large set of i.i.d. random variables (with a “reasonable enough” distribution, *e.g.* a Poisson point process) are described by an *exponential distribution*

$$P_{\text{P}}(s) = e^{-s}, \quad (2.6)$$

which has no level repulsion but rather *level collision*, with $s = 0$ being the most likely value.

A straightforward way to distinguish between ergodic and non-ergodic systems is therefore to compare their level spacing distributions. Indeed, it is known that for many integrable systems, whose eigenvalues can typically be written as sums of independent quantities (think of the transverse-field Ising model, where $E = \sum_k \epsilon_k$ for a set of uncorrelated quasiparticle energies $\{\epsilon_k\}_k$), the spectrum is equivalent to a random collection and displays as a consequence Poisson level statistics.⁹ Likewise, it is possible to show that the eigenvalues of the Anderson model, Eq. (2.1), become uncorrelated in the large-size limit [187]. The same is expected to be true for MBL systems, which possess a hidden integrability (Section 2.2.2).

A particularly expedient characterization of the level spacing spectrum was introduced by Oganesyan and Huse in their aforementioned work [195]. A common problem that one encounters when numerically computing the level spacing distribution $P(s)$ is that it can be prone to large finite-size effect due to its sensitive dependence on the local density of states. This can be alleviated to a large extent by considering a two-point level statistics, the “ r -parameter”:

$$r_n = \frac{\Delta_{n \wedge (n+1)}}{\Delta_{n \vee (n+1)}}, \quad (2.7)$$

with Δ_n and Δ_{n+1} being two consecutive level spacings and $\Delta_{n \wedge (n+1)}$ ($\Delta_{n \vee (n+1)}$) denoting the minimum (maximum) between the two.

⁹The Berry–Tabor conjecture is that this behavior holds true for all quantum systems whose classical counterpart is integrable, provided they do not possess nontrivial symmetry sectors causing commensurabilities in the spectrum [42, 73].




Figure 2.2: Typical behavior of the average level spacing ratio $\langle r \rangle$ for a localizing model — in this case, the 3D Anderson model (2.1), which features a localization transition at the critical value of the disorder strength $W_c \approx 16.5$. In the ergodic phase, $W < W_c$, the r parameter tends to the Wigner–Dyson value $\langle r \rangle \approx 0.53$ as the system size L is increased; in the localized phase, $W > W_c$, it tends instead to the Poisson value $\langle r \rangle \approx 0.39$.

Its distribution can be “surmised” for the Wigner–Dyson case and computed exactly for the Poissonian one [28]:

$$P_{\text{WD}}(r) \approx \frac{27}{4} \frac{r + r^2}{(1 + r + r^2)^{5/2}}, \quad (2.8)$$

$$P_{\text{P}}(r) = \frac{2}{(1 + r)^2}, \quad (2.9)$$

with respective averages $\langle r \rangle_{\text{WD}} = 4 - 2\sqrt{3} \approx 0.5359$ and $\langle r \rangle_{\text{P}} = 2 \log 2 - 1 \approx 0.3863$. A more accurate estimate for the RMT result is obtained by fitting numerical data for large matrices, yielding $\langle r \rangle_{\text{WD}} = 0.5307(1)$ instead. These distributions and averages are rather robust to finite-size effects, which makes them a useful detector for MBL. A typical MBL system will have $\langle r \rangle$ flowing from the Poisson to the Wigner–Dyson value as some control parameter is swept across the localization–delocalization transition (Fig. 2.2).

We mention another possible indicator of non-ergodicity obtained by comparing consecutive eigenstates. If the ETH is satisfied, nearby eigenstates should “look alike”, resulting in a small value of the relative entropy

$$D(p_a \| p_{a+1}) = - \sum_{z=0}^{D-1} p_a(z) \log \frac{p_a(z)}{p_{a+1}(z)} \quad (2.10)$$

between their amplitude distributions $p_a(z) = |\langle z | \psi_a \rangle|^2$ in some fixed computational basis

$\{|z\rangle\}_z$. If $D(p_a||p_{a+1})$ does not remain bounded as the system size is increased, the ETH must be violated.

Entanglement in the eigenstates

Rather than comparing different eigenstates, one can instead look for a violation of the ETH within single eigenstates as well. As mentioned in Section 1.1, the ETH predicts that eigenstates at nonzero energy densities should be similar to random vectors in a Hilbert space. Taken seriously, this prediction implies [196, 226] that their *entanglement entropy*¹⁰ should display a typical volumetric scaling with the subsystem size

$$S_A^{\text{ee}} \sim |A| \quad (\text{ergodic}) \quad (2.11)$$

where $A \amalg A^c$ is any spatial bipartition of the system with $|A| \leq |A^c|$.

For a *bounded* subsystem A , local thermalization (see Definition 1.1.3) requires ρ_A to be equivalent to a Gibbs state $\propto e^{-\beta H_A}$ with inverse temperature dictated by the eigenstate's energy. In this case, it is easy to see that the value of S_A^{ee} should reproduce the thermal entropy:

$$S_A^{\text{ee}} = -\text{Tr}_A \left[\frac{e^{-\beta H_A}}{Z_A} \log \frac{e^{-\beta H_A}}{Z_A} \right] = \beta (\langle H_A \rangle - F_A) = S_A^{\text{th}}. \quad (2.12)$$

As both $\langle H_A \rangle$ and F_A are extensive in the size of A , we recover the volume law (2.11).¹¹

The volumetric scaling of entanglement was indeed found to hold generically in ergodic many-body systems, at least in the bulk of the spectrum.¹² In a lattice of linear size L in d dimensions, the volume law can be written in the form $S_L^{\text{ee}} \sim L^d$, which stays valid in the Hamiltonian limit where the lattice spacing is sent to zero and $|A| \rightarrow \infty$ while L is fixed. Here and later, the subscript L of S_L^{ee} is meant to denote any given subset A whose linear size grows proportionally to L , e.g. one half of the system.

Localized eigenstates, on the other hand, are found to obey an *area law*

$$S_A^{\text{ee}} \sim |\partial A| \quad (\text{localized}), \quad (2.13)$$

with the entropy proportional not to the size of the subsystem, but rather of its boundary. An area law is indicative of short-range correlations: roughly speaking, if one conceives of

¹⁰The bipartite entanglement entropy of state ρ relative to the Hilbert space decomposition $\mathcal{H} = \mathcal{H}_A \oplus \mathcal{H}_B$ is defined as $S_A^{\text{ee}} = -\text{Tr}_A[\rho_A \log \rho_A] = S_B^{\text{ee}}$, with $\text{Tr}_A[\cdot]$ the partial trace over a basis of \mathcal{H}_A and $\rho_A = \text{Tr}_B[\rho]$. A spatial bipartition of a system $\Omega = A \amalg B$ induces a natural bipartition of the Hilbert space when \mathcal{H} has a tensor product structure.

¹¹Notice that the relation (2.12) does not hold for the *unbounded* set A^c , preventing us from erroneously concluding that $S_{A^c}^{\text{ee}} = S_{A^c}^{\text{th}} \propto |A^c|$, which would be at odds with the mirror property of the entanglement entropy $S_A^{\text{ee}} = S_{A^c}^{\text{ee}}$.

¹²We gloss over the recent topic of quantum scars [253], whereby a quantum system may contain a small ($O(n)$) number of anomalous, low-entanglement states in the middle of the spectrum. The physics behind this form of weak ergodicity breaking is still unclear, and the matter of exciting research.

S_A^{ee} as a quantifier for correlations across the border between A and its complementary, for this number to be proportional to $|\partial A|$ it is necessary that only degrees of freedom within a bounded distance from the border are correlated, unlike the volumetric case where correlations span the entire system. The Hamiltonian-limit-friendly version of Eq. (2.13) is $S_L^{\text{ee}} \sim L^{d-1}$. The area law for MBL eigenstates is entirely analogous to the one obeyed by ground states of local Hamiltonians [61, 118, 238] (though gapless systems in 1D display a logarithmic anomaly [149]).

The properties of the *entanglement spectrum* in MBL systems have been discussed as well. This is defined as the set of eigenvalues $\{\lambda_s\}_s$ of the “entanglement Hamiltonian” $H_{\text{ent}} = -\log \rho_A$, and as such contains at least as much information as the entanglement entropy alone, $S = \sum_s e^{-\lambda_s} \lambda_s$. The distribution $P(\lambda)$ of these eigenvalues (also called the “entanglement density of states”) was found to present a markedly different structure in localized and ergodic models [112, 113].

Quench dynamics of entanglement

Even though entanglement is area-law bounded in MBL eigenstates, quantum information still *can* travel across MBL systems, albeit very slowly.

Entanglement dynamics in a quantum system is usually studied in a prototypical *quantum quench* setup: the system is prepared in a certain initial state $|\psi_0\rangle$, and then suddenly let evolve with the evolution operator $U(t) = e^{-itH}$, with H time-independent. The initial state is usually required to be simple to describe/prepare, for instance being the ground state of some known Hamiltonian H_0 . The quench protocol then coincides with the nonadiabatic transformation

$$H(t) = \begin{cases} H_0 & t < 0 \\ H & t \geq 0 \end{cases}. \quad (2.14)$$

Obviously one wants $[H_0, H] \neq 0$ to induce a nontrivial time evolution. We also assume that the initial state is a product state, so that the initial entanglement is zero. This framework is used in experimental, numerical and analytical settings alike, though in the case of entanglement entropy, which is not observable, the first case can be ruled out.

In ergodic systems, the spreading of entanglement happens quickly (probably ballistically in 1D [143]), and the entanglement entropy saturates to a thermal value $S_A^{\text{ee}} \rightarrow s(E_0)|A|$, with $s(E_0)$ the microcanonical entropy density. This is to be expected (at least for bounded A), as the final state is locally thermal and each degree of freedom can be in any of $e^{s(E_0)}$ different microstates.

On the contrary, it was shown [34, 261] that the half-system entanglement entropy following a quench in the MBL phase in one dimension exhibits a slow growth

$$S_L^{\text{ee}}[\psi(t)] \sim \log t, \quad (2.15)$$

(after a short period of rapid propagation of information within the localization length) which eventually saturates (in finite systems) to an extensive, but *subthermal* value, $S_L^{\text{ee}} \rightarrow cL/2$, with $c < s(E_0)$. This saturation can be interpreted as a form of incomplete thermalization, where additional conservation laws other than the conservation of energy force the (local) steady state to be a generalized Gibbs ensemble [216] with multiple conserved charges $\{Q_j\}_{j=0,1,\dots'}$

$$|\psi(t)\rangle \langle \psi(t)| \longrightarrow \rho^{\text{GGE}} \propto e^{-\sum_j \theta_j Q_j} \quad (2.16)$$

with $\theta_0 = \beta$, $Q_0 = H$, $[Q_j, H] = 0$ for all j and the values of θ_j determined by the initial condition.

Absence of transport

The historically prominent feature of many-body localization, as originally defined by Basko *et al.*, is the suppression of quantum transport of conserved charges through the system, in analogy with the phenomenon of AL in noninteracting systems.

Suppose that Q is a conserved local operator, that is to say,

$$Q = \sum_i q_i \quad (2.17)$$

for some local density q_i . Familiar examples would include the energy H and the total magnetization S^z . For such Q , one can define the corresponding current J from the continuity equation, which will itself be a local operator, $J = \frac{1}{\mathcal{V}} \sum_i j_i$, with \mathcal{V} the system volume.

From linear response theory one then has a corresponding conductivity in terms of the current–current autocorrelator, as expressed by the Kubo formula [165,219]

$$\sigma(\omega) = \frac{\mathcal{V}}{\omega} \Re \int_0^\infty dt \langle [J(t), J(0)] \rangle e^{-i\omega t}, \quad (2.18)$$

where $\langle \cdot \rangle$ is the thermal expectation, and only the real (dissipative) part of the conductivity is considered. Eq. (2.18) be converted to a Lehmann representation by working in the energy basis $\{|a\rangle\}_a$, which gives

$$\sigma(\omega) = \pi \mathcal{V} \sum_{ab} \frac{e^{-\beta E_b}}{Z} \frac{1 - e^{-\omega\beta}}{\omega} |\langle a|J|b\rangle|^2 \delta_\eta(\omega - E_a + E_b). \quad (2.19)$$

Here δ_η represents a Cauchy distribution of small width η , used to smoothen out the density of states at finite size (η should be taken to 0 *after* the thermodynamic limit). The DC conductivity is then defined as the $\omega \rightarrow 0$ limit of the above formula, namely

$$\sigma_{\text{dc}} = \pi \beta \mathcal{V} \sum_{ab} \frac{e^{-\beta E_b}}{Z} |\langle a|J|b\rangle|^2 \delta_\eta(E_a - E_b). \quad (2.20)$$

	Ergodic	MBL
Level spacing distribution	Wigner–Dyson (level repulsion)	Poissonian (level collision)
Average r -parameter	≈ 0.5307	≈ 0.3863
Eigenstate entanglement entropy	volume law	area law
Post-quench entanglement entropy	quick $O(t)$ growth up to thermal value	slow $O(\log t)$ growth up to subthermal value
DC conductivity	nonzero	zero

Table 2.1: Summary of the main phenomenological differences between ergodic and MBL systems.

Just like in AL systems, many-body localization causes this coefficient to exactly vanish, $\sigma_{dc} = 0$, even at nonzero temperatures (see the next section for a proof in terms of local integrals of motion).

Although this was the original characterization of MBL, it is nowadays not considered to be the best way to actually define the localized phase. For instance, as pointed out in Ref. [192], a Floquet system without any local conserved charge Q — and hence no notion of DC transport — may still undergo a recognizable MBL phase transition. As we will see in the next section, a general definition of MBL can be given, at least for the case of fully localized spectrum, in terms of the existence of *quasilocal* conserved operators.

Summary

The observations of this section are summarized in Table 2.1.

2.2.2 The LIOM theory of MBL

The theoretical understanding of many-body localization has been much improved as a consequence of its formulation, by now canonical, as a kind of emergent integrability, described by the framework of *(quasi)local integrals of motion* (LIOMs).

The idea behind LIOMs is that in a localized system, local degrees of freedom should be only weakly dressed by the interactions compared to the noninteracting regime [125, 228]. Concretely, consider a noninteracting system of qudits with local dimension q (e.g. $q = 2$ for spins- $\frac{1}{2}$). Upon diagonalization, the system’s Hamiltonian will be expressible

in the form $H = \sum_i \epsilon_i \sigma_i^z$ up to a constant shift, where σ_i denotes a diagonal operator labeling the state of the system in some computational basis (this can also be thought of as an “occupation number”, where we allow up to $q - 1$ particles per site). The set of corresponding eigenvalues, $\{\epsilon_i^z\}_i$, uniquely identifies the system’s eigenstates.

Now, if we introduce a slight perturbation of these degrees of freedom resulting from local interactions, the mutually commuting operators $\{\tau_i^z\}_i$ determining the new good quantum numbers for our system will be expressed generically as sums of multiary non-diagonal terms in the σ_i^α ’s ($\alpha \in \{x, y, z\}$):

$$\tau_i^z = \sigma_i^z + \sum_{j_1 j_2} \sum_{\alpha_1 \alpha_2} \theta_{ij_1 j_2}^{\alpha_1 \alpha_2} \sigma_{j_1}^{\alpha_1} \sigma_{j_2}^{\alpha_2} + \sum_{j_1 j_2 j_3} \sum_{\alpha_1 \alpha_2 \alpha_3} \theta_{ij_1 j_2 j_3}^{\alpha_1 \alpha_2 \alpha_3} \sigma_{j_1}^{\alpha_1} \sigma_{j_2}^{\alpha_2} \sigma_{j_3}^{\alpha_3} + \dots \quad (2.21)$$

with the couplings $\{\theta_{ij_1 \dots j_r}^{\alpha_1 \dots \alpha_r}\}_{r \geq 2; j_1, \dots, j_r \in [n]}$ exponentially bounded in the spatial separations:

$$\theta_{ij_1 \dots j_r}^{\alpha_1 \dots \alpha_r} \leq e^{-\frac{\max_k |i-j_k|}{\xi_1}} \quad (2.22)$$

for some *localization length* ξ_1 .

A unitary change of basis (2.21) satisfying Eq. (2.22) is called *quasilocal*, and since the $\{\tau_i^z\}_i$ are by definition conserved, they are referred to as (quasi)local integrals of motion (LIOMs). One usually has in mind the $q = 2$ case, where LIOMs are also called *l-bits* (as in “localized bits”). Indeed, notice that despite the complicated form of Eq. (2.21), unitarity implies that the l-bits must have eigenvalues ± 1 . The latter terminology has the advantage of explicitly reminding one that the τ_i^z operators are constructed out of physical “bits” by means of a quasilocal change of basis. In fact, not all integrals of motion that happen to be quasilocal can be obtained this way.¹³

Since all LIOMs commute with the Hamiltonian, this has to be diagonal in them:

$$H = \sum_{j_1} J_{j_1} \tau_{j_1}^z + \sum_{j_1 j_2} J_{j_1 j_2} \tau_{j_1}^z \tau_{j_2}^z + \sum_{j_1 j_2 j_3} J_{j_1 j_2 j_3} \tau_{j_1}^z \tau_{j_2}^z \tau_{j_3}^z + \dots \quad (2.23)$$

for some set of $\{J_{j_1 \dots j_r}\}_{r \geq 1; j_1, \dots, j_r \in [n]}$ which must also be exponentially localized in real space:

$$J_{j_1 \dots j_r} \leq e^{-\frac{\max_{k,l} |j_k - j_l|}{\xi}}. \quad (2.24)$$

The existence of LIOMs was rigorously proved for a class of spin chains in one dimensions under physically reasonable assumptions [128], and their construction was carried out perturbatively in the interaction strength [219].

The LIOM formalism is a powerful explanatory tool for almost all the properties of the MBL phase discussed in the previous section, and more. For instance, since the eigenstates

¹³See Ref. [66] for an explicit construction of local integrals of motion which are not l-bits.

of (2.23) only differ by a quasilocal transformation from eigenstates of the $\sum_i \sigma_i^z$ operator, which are spatially factorized, the only sites contributing to their entanglement entropy are those within distance $O(\xi_1)$ from the boundary of the bipartition, resulting in an area law.

Another example is the vanishing of the dc conductivity σ_{dc} , mentioned in Section 2.2.1.¹⁴ Consider the case of a local current with density j_i , and suppose that each j_i is supported on a bounded set of diameter ζ . Assume that a set of LIOMs exists, and take two eigenstates $|a\rangle$ and $|b\rangle$ such that $\tau^z |a\rangle = \tau_a |a\rangle$, $\tau^z |b\rangle = \tau_b |b\rangle$ for some distinguished l -bit τ , with $\tau_a \neq \tau_b$. Then, if $r \gg \zeta$ we have that either $[j_i, \tau] = 0$ or $[j_{i+r}, \tau] = 0$ due to the locality of the involved operators. Therefore, the amplitudes

$$\langle a|j_i|b\rangle = \frac{\langle a|[j_i, \tau]|b\rangle}{\tau_b - \tau_a} \quad \text{and} \quad \langle b|j_{i+r}|a\rangle = \frac{\langle b|[j_{i+r}, \tau]|a\rangle}{\tau_a - \tau_b} \quad (2.25)$$

cannot both be nonzero. Therefore, when we rewrite Eq. (2.20) for our local current $J = \frac{1}{\mathcal{V}} \sum_i j_i$,

$$\sigma_{\text{dc}} = \frac{\pi\beta}{\mathcal{V}} \sum_{ab} \sum_{ir} \frac{e^{-\beta E_b}}{Z} \langle a|j_i|b\rangle \langle b|j_{i+r}|a\rangle \delta(E_a - E_b), \quad (2.26)$$

we can see that the sum over r is actually restricted to $r \lesssim \zeta$. Moreover, since the support of j_i can only contain a bounded number of LIOMs (roughly $O(c^{\zeta^d})$ in d dimensions, for some c), the state $j_i |a\rangle$ can only overlap nontrivially with a bounded number of eigenstates $|b\rangle$, because if $|a\rangle$ and $|b\rangle$ differ by LIOMs supported outside of the support of j_i , the operator cannot “fix” the discrepancy and $\langle b|j_i|a\rangle = 0$ as a consequence. Therefore, the sum over b is restricted as well. As a result, the number of energy levels contained in the sum becomes a negligible fraction of the spectrum and the delta approximant δ_η — and hence the conductivity — will vanish with probability one when taking the $n \rightarrow \infty$ and $\eta \rightarrow 0$ limits.¹⁵

A full discussion of the LIOM picture and its implications would take us far afield, so we refer to Ref. [129] for a more complete discussion. For the rest of the chapter, we are going to have to abandon (at least in part) the canonical understanding of many-body localization exposed so far, as we explore the possibility of witnessing yet different phenomena in the vast ocean of disordered interacting models.

2.3 Isotropic disordered models

We now introduce and motivate the principal topic of concern for the present chapter, namely, the interplay of *symmetry and localization*. As we shall see, a certain clash arises

¹⁴The following argument is adapted from Ref. [219].

¹⁵The $r = 0$ term, which is never discarded by the delta function, does not contribute as $\langle a|j_i|a\rangle = 0$ for time-reversal-invariant models.

between the two notions whenever the symmetry in question is non-Abelian, as the non-trivial structure of its faithful representations tries to impose additional constraints to the quantum states on top of the LIOM structure. The two sets of constraints turn out to be incompatible, and a compromise must be reached.

In this section, we take a detailed look at the relatively simple and physically relevant $SU(2)$ symmetry. Many constructions and arguments presented for this case, in particular regarding the “fusion” of local degrees of freedom into nonlocal multiplets, can in principle be generalized to more complex Lie groups such as $SU(N \geq 3)$, whose representation theory is “only” complicated by the appearance of multiple Dynkin labels rather than just the total spin.

The result of the aforementioned symmetry–localization competition is yet to be uncovered in its full generality. In this view, Section 2.4 is a first step in the direction of gathering insightful information in an exemplary model, including a proposal for the physical mechanism underlying the observed phenomenology.

2.3.1 Motivation: the Hubbard model with local disorder

The quantum Heisenberg model is described by the Hamiltonian

$$H = -J \sum_{\langle ij \rangle} \mathbf{s}_i \cdot \mathbf{s}_j, \quad (2.27)$$

where $s_i^\alpha = \sigma_i^\alpha/2$ is the spin- $\frac{1}{2}$ operator acting on site i , and we use the vector notation $\mathbf{s}_i = (s_i^x, s_i^y, s_i^z)$ for the internal degrees of freedom. This Hamiltonian is characterized by an isotropic two-body interaction, motivating the alternative name ‘XXX’ for the Heisenberg model in contrast to its partially or fully anisotropic generalizations, the XXZ and XYZ models. The sum is carried out over first-neighbor sites in some prescribed interaction graph, determining the dimensionality of the system.

As the name suggests, the Heisenberg Hamiltonian has been used as a description of (anti)ferromagnetic materials since the early days of quantum mechanics [121]. The spin variables are meant to represent the effective magnetic moments of the atoms interacting with each other through a short-range interaction. The exact nature of the interaction itself is left unspecified, only retaining the assumption that whatever is happening at a fundamental level, the interaction energy can ultimately be considered, to a satisfactory degree of approximation, to only depend on the dipoles’ mutual orientation. This is in fact what one *means* by “magnetic interaction” in this context. As a matter of fact, the actual magnetic dipoles of the atoms are not the only — nor the main — underlying source of magnetic interactions in typical (anti)ferromagnets. Other mechanisms include spin–orbit coupling and, especially, the Coulomb interaction between shallow electrons (in turn classified into direct exchange, superexchange, *etc.*) [27].

Let us briefly discuss the Heisenberg Hamiltonian in its one-dimensional version. This model was provided an exact solution by Bethe [43], whose technique was later harnessed

to great effect for tackling other 1D statistical models [37]. The constant J , the only parameter of the fully isotropic model, determines the energy units and can be set to 1 as far as the spectrum is concerned. Its sign, however, determines the low-energy physics of the system. If $J > 0$, the system is ferromagnetic: dipoles tend to align at low temperature so as to maximize the total spin, as can be seen by rewriting the Hamiltonian in the form

$$H = -J \sum_i \left[(\mathbf{s}_i + \mathbf{s}_{i+1})^2 - \frac{3}{2} \right], \quad (2.28)$$

which implies that H is minimized by any state which is obtained by recursively combining all pairs of neighboring spins into maximal-spin multiplets, thereby creating a maximal-spin state in turn.¹⁶ If the number of spins is n , the maximal spin is $n/2$ and there are $n + 1$ ground states. This degeneracy reflects the high degree of symmetry of H , as we are free to change magnetization sector without affecting the energy. Excitations over a ground state are created by lowering the value of the total spin by 1, creating a quasiparticle (*magnon*).

Taking $J < 0$ gives instead the antiferromagnetic Heisenberg model, which can be shown to have a unique ground state living in the $M = 0$ eigenspace of the total magnetization. This can be thought of as a sea of free magnons, and elementary excitations are now understood as creating a different kind of quasiparticle (a *spinon* pair). At higher energy, bound states of magnons can also be created. [108]

As is typical for one-dimensional spin models, an alternative interpretation of the Heisenberg model is obtained by fermionizing the Hamiltonian via a Jordan–Wigner transformation. The procedure is described in detail in Appendix A. Upon substituting

$$s_i^+ = s_i^x + i s_i^y = e^{i\pi \sum_{j<i} c_j^\dagger c_j} c_i \quad s_i^z = \frac{1}{2} - c_i^\dagger c_i,$$

with the $\{c_i\}_i$ operators satisfying the canonical anticommutation relations $\{c_i, c_j^\dagger\} = \delta_{ij}$ and $\{c_i, c_j\} = 0$, the Heisenberg Hamiltonian is revealed to be equivalent to a model of interacting fermions:

$$H = -J \sum_i \left[\frac{1}{2} (c_i^\dagger c_{i+1} + c_{i+1}^\dagger c_i) + \left(c_i^\dagger c_i - \frac{1}{2} \right) \left(c_{i+1}^\dagger c_{i+1} - \frac{1}{2} \right) \right], \quad (2.29)$$

where for simplicity we quote the form valid for open boundary conditions. In this Hamiltonian, one can recognize a nearest-neighbor hopping term

$$-\frac{J}{2} \sum_i (c_i^\dagger c_{i+1} + c_{i+1}^\dagger c_i),$$

¹⁶This combination or “fusion” procedure will be fully explained in Section 2.3.3.

a two-body interaction term, also nearest-neighbor,

$$-J \sum_i c_i^\dagger c_i c_{i+1}^\dagger c_{i+1}$$

and an extra on-site potential which at this stage merely provides an energy shift if one restricts oneself to a fixed (conserved) total number of fermions.

A glassy version of Eq. (2.27) is obtained by treating the spin–spin couplings as quenched random variables, resulting in the *random-bond Heisenberg model* (RBHM)

$$H = - \sum_{ij} J_{ij} \mathbf{s}_i \cdot \mathbf{s}_j. \quad (2.30)$$

In view of Eq. (2.29), the one-dimensional version of this model can be interpreted as a fermionic model with quenched disorder, though unlike typical case studies for disordered fermionic systems, the randomness enters the kinetic as well as the interaction term, making the dynamics nontrivial even in the absence of interaction.¹⁷

Indeed, one can do the opposite and start from a more familiar fermionic system to motivate the study of Eq. (2.30). Consider a chain inhabited by two distinct, but identically-behaving, fermionic species, canonically anticommuting with each other, where fermions can hop between neighboring sites and interact when sitting on the same site. This is the Hubbard model [92, 123] describing correlated electrons in a solid, where a single band is considered and the Coulomb interaction is simplified down to a contact interaction.¹⁸ Unlike the familiar Hubbard model, we consider a disordered version of the system with random on-site energies (*i.e.* a site-dependent chemical potential) $\{\mu_i\}_i$:

$$H = -t \sum_{s=\uparrow,\downarrow} \sum_i \left(c_{i,s}^\dagger c_{i+1,s} + c_{i+1,s}^\dagger c_{i,s} \right) + U \sum_i n_{i\uparrow} n_{i\downarrow} + \sum_i \mu_i (n_{i\uparrow} + n_{i\downarrow}), \quad (2.32)$$

¹⁷In a single-particle system it is possible to compute the energy-dependent Lyapunov exponent $\lambda(E)$ in terms of the density of state $\rho(E)$ via the Herbert–Jones–Thouless formula [122, 250]

$$\lambda(E) = \int \rho(E') \log|E - E'| dE' + \text{const}. \quad (2.31)$$

One then has $\langle x|E \rangle \sim e^{-\lambda(E)x}$ for an eigenstate of energy E , which means that $1/\lambda(E)$ is a localization length. In the random hopping model, one finds a nonzero $\lambda(E)$ (*i.e.* Anderson localization) through the entire spectrum, except for $\lambda(E \rightarrow 0) \rightarrow 0$. This signals the presence of arbitrarily wide eigenstates at low energy (*i.e.* there is no dynamical localization, in the language of Section 2.1.1). There is a nice correspondence between a chain of noninteracting quantum particles and a classical system of coupled strings. The emergence of a delocalized state at $E \rightarrow 0$ corresponds to the presence of a massless (Goldstone) zero mode in the classical system.

¹⁸Although this was not the case in the 1960s, it is nowadays possible to simulate the Hubbard model in cold atom systems [93, 139]. These provide a controlled, tunable and upscaled environment which is ideal to test the rich predictions of the model, such as the metal–insulator transition or the BEC–BCS crossover. In these cold atom experiments, the Hubbard Hamiltonian is typically realized using two hyperfine levels of alkaline atoms such as ${}^6\text{Li}$ or ${}^{40}\text{K}$ trapped in a laser array.

where $c_{i\uparrow}^\dagger$ and $c_{i\downarrow}^\dagger$ create the two kinds of fermion at site i and $n_{i,s} = c_{i,s}^\dagger c_{i,s}$. Notice that H separately conserves both fermionic numbers N_\uparrow, N_\downarrow , *i.e.* there is a $U(1) \otimes U(1)$ symmetry in place. In fact, the Hubbard model has a much larger symmetry group, which most notably includes the non-Abelian “flavor” $SU(2)$ generated by the operators

$$S^\alpha = \sum_i S_i^\alpha, \quad \alpha \in \{x, y, z\}, \quad (2.33)$$

where

$$S_i^\alpha = \frac{1}{2} \sum_{s,s'} c_{i,s}^\dagger (\sigma^\alpha)_{s,s'} c_{i,s'} \quad (2.34)$$

($\{\sigma^\alpha\}_{\alpha \in \{x,y,z\}}$ being the Pauli matrices). One can verify that S^x, S^y and S^z form a Pauli algebra.

By performing the inverse of the Jordan–Wigner transformation quoted above¹⁹, we can make Eq. (2.32) into a spin model of two XX chains with parallel ZZ couplings and local longitudinal fields:

$$\begin{aligned} H = & -t \sum_i (\zeta_i^+ \zeta_{i+1}^- + \zeta_{i+1}^+ \zeta_i^- + \tau_i^+ \tau_{i+1}^- + \tau_{i+1}^+ \tau_i^-) \\ & + U \sum_i \zeta_i^z \tau_i^z + \sum_i \epsilon_i (\zeta_i^z + \tau_i^z) \end{aligned} \quad (2.35)$$

(up to a constant shift²⁰), with the spin operators²¹ ζ_i^a and τ_i^a mutually commuting and $\epsilon_i = \mu_i + U/2$. At this point, one recognizes that the interaction part (line two in Eq. (2.35)) is quadratic in $\zeta_i^z + \tau_i^z$. Namely, we have

$$H_{\text{inter}} = H(t = 0) = \frac{U}{2} \sum_i C_i^2 + \sum_i \epsilon_i C_i \quad (2.36)$$

(again, constants aside), with $C_i = \zeta_i^z + \tau_i^z = n_{i\uparrow} + n_{i\downarrow} - 1 = n_i - 1$, a shifted version of the operator n_i counting the total number of fermions currently present on site i regardless

¹⁹One needs to be careful in this case to properly define the transformation so as to account for both fermionic species. The string operators in particular need to be defined correctly. A convenient way to do so is to imagine that $c_{i,s}^\dagger$ creates a (spinless) fermion at site (i, s) , and order the resulting $2L$ sites such that $(1, \uparrow) < \dots < (L, \uparrow) < (1, \downarrow) < \dots < (L, \downarrow)$.

²⁰One also needs to assume that N_\uparrow and N_\downarrow are both odd, if periodic boundary conditions are enforced. More generally, one needs to replace the $i = L$ term of the sum in the kinetic part with

$$\Delta H_{\text{PBC}} = +t \left[(-1)^{N_\uparrow} \zeta_L^+ \zeta_1^- + (-1)^{N_\downarrow} \tau_L^+ \tau_1^- + \text{H.c.} \right].$$

²¹We employ the ζ notation to prevent confusion with the standard Pauli matrices, which differ from spin operators by a factor of two.

of their species. Operator C_i can be referred to as “charge”, because it measures signed deviations from a state with exactly one fermion per site, the lack of a fermion being interpreted as a (-1) -charged “holon” and the excess of a fermion as a $(+1)$ -charged “doublon”.

In the assumption that the total number of particles equals the number of sites, *i.e.* assuming half-filling, the energetic cost for two opposite-spin electrons to sit on the same site is U plus a correction from the site-dependent chemical potential. In the clean model, the correction is zero and creating a doublon is always disadvantageous at $t/U = 0$. Therefore, the ground manifold \mathcal{H}_0 is comprised of the 2^L degenerate states of uniform charge $\{C_i = 0, \forall i\}$. This is used as the basis of a perturbation theory in the kinetic term. We can do the same here as long as we assume that the local fields $\{\epsilon_i\}_i$ are contained inside some bandwidth W which is much smaller than U . The subspace \mathcal{H}_0 is no more, strictly speaking, a ground manifold, but it is still separated from the rest of the Hilbert space by an insurmountably large energy barrier, and we are justified in considering the dynamics confined to this subspace.

Out of the four possible states of the i -th site Hilbert space,

$$|\emptyset_i\rangle, |\uparrow_i\rangle, |\downarrow_i\rangle, |(\uparrow\downarrow)_i\rangle,$$

only $|\uparrow_i\rangle$ and $|\downarrow_i\rangle$ are retained to span the desired subspace. Here $|\uparrow_i\rangle$ represents a spin-up fermion at site i , *i.e.* $|\uparrow_i\rangle = c_{i\uparrow}^\dagger |\emptyset_i\rangle$, and similarly for the other symbols. In spin language, these are equivalent — up to unimportant signs from the Jordan–Wigner transformation — to

$$|(++)_i\rangle, |(+ -)_i\rangle, |(- +)_i\rangle, |(--)_i\rangle,$$

where *e.g.* the second state is a simultaneous eigenstate of ζ_i^z and τ_i^z with respective eigenvalues $+\frac{1}{2}$ and $-\frac{1}{2}$.

On half-filled states, the interaction term Eq. (2.36) vanishes identically and the kinetic term, as projected onto \mathcal{H}_0 , acts as an $O(t^2/U^2)$ hopping term by first displacing a fermion to a neighboring site — creating a virtual holon–doublon pair — and later restoring the zero-charge condition by moving the other fermion (necessarily of opposite spin, as the virtual transition would otherwise be Pauli-prohibited) to the original site:

$$\begin{array}{ccc}
 & & |\emptyset_i (\uparrow\downarrow)_{i+1}\rangle \\
 & \nearrow & \searrow \\
 |\uparrow_i \downarrow_{i+1}\rangle & & \\
 & \searrow & \nearrow \\
 & & |(\uparrow\downarrow)_i \emptyset_{i+1}\rangle \\
 & & \searrow \nearrow \\
 & & |\downarrow_i \uparrow_{i+1}\rangle.
 \end{array} \tag{2.37}$$

Performing the standard Rayleigh–Schrödinger perturbation theory reveals that the effective perturbation

$$H_{\text{eff}}^{(2)} = V \frac{1}{E_0 - H_0} \Big|_{\mathcal{H}_0^\perp} V \tag{2.38}$$




Figure 2.3: The distribution of random couplings $\{J_i\}_i$ in the effective Heisenberg Hamiltonian, Eq. (2.39), stemming from a uniform distribution of local fields $\{\epsilon_i\}_i$ in the Hubbard model, Eq. (2.32).

(with H_0 and V the interaction and kinetic parts of the Hamiltonian, respectively, and $E_0 = \sum_i \epsilon_i$ the unperturbed ground-state energy) is indeed equivalent to a random-bond Heisenberg model (up to a constant shift) acting on fictitious spins- $\frac{1}{2}$:

$$H_{\text{eff}}^{(2)} = \sum_i J_i \mathbf{s}_i \cdot \mathbf{s}_{i+1}, \quad (2.39)$$

where

$$J_i = \frac{2t^2U}{U^2 - (\epsilon_i - \epsilon_{i+1})^2} \quad (2.40)$$

is the effective coupling computed from the aggregate second-order transition amplitude of processes Eq. (2.37) (namely the kinetic exchange [20]) and $\{\mathbf{s}_i\}_i$ are spin-like operators defined by their action on \mathcal{H}_0 :

$$2s_i^z \begin{array}{c} \curvearrowright \\ \curvearrowleft \end{array} |(+ -)_i\rangle \begin{array}{c} \xrightarrow{s_i^-} \\ \xleftarrow{s_i^+} \end{array} |(- +)_i\rangle \begin{array}{c} \curvearrowleft \\ \curvearrowright \end{array} -2s_i^z.$$

The Hamiltonian (2.39) describes therefore the dynamics of L fictitious spins interacting isotropically with their neighbors. Adopting the Eq. (2.29) view, they can also be thought of as fermions hopping on a chain through a nonuniform hopping matrix and experiencing a disordered two- and one-body potential (the disorder in the interaction at one site being perfectly correlated to the one in the hopping about the same site, as they are both specified by J_i). Notice how the global SU(2) symmetry of the Hubbard model survives in $H_{\text{eff}}^{(2)}$.

A few words are in order concerning the distribution of the effective couplings J_i . First off, notice from Eq. (2.40) that all couplings are positive, corresponding to an *antiferromagnetic* model. The second observation is that the distribution of J_i is induced from that of

$\{\epsilon_i\}_i$ in a nontrivial way. For instance, if the ϵ_i are uniformly distributed in $[-W/2, W/2]$, the marginal²² distribution for the J_i couplings is (Fig. 2.3)

$$P(J) = \frac{t^2 U}{W J^2} \left[\left(1 - \frac{t^2}{U J} \right)^{-1/2} - \frac{U}{W} \right] \quad \frac{t^2}{U} < J \leq \frac{t^2 U}{U^2 - W^2}, \quad (2.41)$$

which diverges with exponent $-1/2$ at the minimum possible coupling $J \rightarrow \frac{t^2}{U}$, corresponding to the “resonant” condition $\epsilon_i \approx \epsilon_{i+1}$, and decays monotonously until the maximum possible coupling $J = \frac{t^2 U}{U^2 - W^2}$, corresponding to $|\epsilon_i - \epsilon_{i+1}| = W$.

As we will see shortly, the most relevant feature of this distribution is that it is singular and with compact support. In Section 2.4, we will consider a simpler pdf with the same relevant features. One can therefore view our following numerical analysis as having in mind this uniform-disorder, half-filled Hubbard model as a reference point.

2.3.2 Low-temperature thermodynamics

The antiferromagnetic Hamiltonian (2.30) was first considered by Bulaevskii *et al.* [59] as a model for certain organic molecules, such as complex salts of tetracyanoquinodimethane (TCNQ), whose spatial structure can be regarded as a one-dimensional electronic channel. In this work, the authors assumed that the fermionic counterpart of the model could be described as a Fermi–Landau liquid with a singular density of states $\rho(\epsilon)$ at energy density $\epsilon \rightarrow 0$. The assumption allowed them to explain the magnetic phenomenology observed in TCNQ complexes at low temperature, in particular the singular $\chi(T) \sim T^{-\alpha}$ dependence ($0 < \alpha < 1$) of the magnetic susceptibility.

Bulaevskii *et al.*’s assumptions admittedly rested on rather shaky grounds and were later challenged by Theodorou and Cohen [247], who went on to propose an alternative approach. However, it was not until Dasgupta and Ma [74, 164] entered the stage that the problem was seen to a satisfactory end. They devised a form of spin decimation, later named strong-disorder renormalization group (SDRG), that enabled them to study the low-energy properties of the RBHM by hierarchically simplifying the problem’s interaction graph. As is typically the case in decimation algorithms, the relevant feature of this simplification is that it reduced the number of degrees of freedom while preserving the form of the Hamiltonian, allowing the authors to define a flow in model space as parametrized by the underlying coupling distribution $P(J)$. The fact that the flow admits an attractor implies the presence of universal low-energy properties.

The SDRG step proceeds as follows. Suppose that the coupling constants J_i are independent and distributed according to a compactly-supported pdf $P_0(J)$, $0 \leq J \leq \Omega_0$. We are considering a system in the thermodynamic limit, so that $P_0(J)$ is not simply the

²²Notice that as both J_i and J_{i+1} depend on ϵ_{i+1} , the correlation between them is nonzero, so joint distributions are not products of their marginals.

underlying distribution of the realization of couplings $\{J_i\}_i$, but in fact coincides with the *empirical* distribution of couplings. In particular, we have $\Omega_0 = \max_i J_i$. Now focus on one pair of neighboring spins, say 1 and 2, whose coupling has the maximal value, $J_1 = \Omega_0$. Due to the independence assumption, it will almost always be the case that the couplings, J_0 and J_2 immediately to the left and right of the pair will be smaller than J_2 . Suppose in fact that they are considerably smaller,

$$J_0, J_2 \ll J_1 = \Omega_0, \quad (2.42)$$

which will typically be the case if $P_0(J)$ is biased toward small J , as in Fig. 2.3. Then, it is possible to treat the four-spin system $\{\mathbf{s}_0, \mathbf{s}_1, \mathbf{s}_2, \mathbf{s}_3\}$ perturbatively in $J_{0\vee 2}/\Omega_0$, where $J_{0\vee 2} = \max\{J_0, J_2\}$, by considering the 1–2 interaction as an (easily solved) unperturbed system and the “wing” interactions 0–1 and 2–3 as perturbations. Expressing the ground-state energy of the four-spin system to second order in perturbation theory, one finds that the low-energy properties of H are encoded into an effective Hamiltonian H_{eff} with one fewer pair of spins, in the sense that H and H_{eff} share the same low-energy spectrum up to order $O(J_{0\vee 2}^2/\Omega_0^2)$. This amounts to a decimation step where spins 1 and 2 are frozen in a singlet — thereby contributing only a constant term in H — while spins 0 and 3 are directly linked via a new coupling

$$J'_0 = \frac{J_0 J_2}{2\Omega_0}. \quad (2.43)$$

Notice that J'_0 is (much) less than both $J_0/2$ and $J_2/2$, further biasing the coupling distribution toward small J . Moreover, after performing this decimation for all maximally-coupled pairs, the value Ω_0 is no longer allowed for any of the remaining couplings, shifting the maximum down to a new value $\Omega < \Omega_0$.

The precise change in the distribution of couplings is described by the following equation:

$$x \frac{\partial P}{\partial x} - \Omega \frac{\partial P}{\partial \Omega} = P(1, \Omega) \Omega^2 \int_0^1 dx_1 dx_2 P(x_1, \Omega) P(x_2, \Omega) \delta\left(x - \frac{x_1 x_2}{2}\right), \quad (2.44)$$

where Ω is the “running maximum” of the distribution’s support, $x = J/\Omega$ the value of the coupling normalized in $[0, 1]$ and $P(x, \Omega)$ the distribution of normalized couplings when enough decimation has taken place so that the current maximum coupling is Ω . The initial condition is obviously $P(x, \Omega_0) = P_0(\Omega_0 x)$. Notice that Ω can be rightfully interpreted as the global energy scale of the system. The SDRG procedure is then “integrating out” the high-energy degrees of freedom similarly to the ordinary Wilson renormalization group.

As previously mentioned, the SDRG procedure is expected to provide a good approximation of the low-energy physics as long as the initial pdf $P_0(J)$ is not too skewed toward large couplings, which would make assumption (2.42) collapse. The good news is that once the decimation has started, the situation can only get better as the newly introduced couplings are small by construction. Moreover, Dasgupta and Ma mention that even in the

clean model $P_0(J) = \delta(J - \Omega_0)$, the value of the ground-state energy computed from SDRG agrees to within $\sim 2\%$ with the exact result. One expects the disordered case to always behave better than the clean one, as the latter maximally violates assumption (2.42).

Although Eq. (2.44) cannot be solved exactly, an approximate analytical treatment as well as numerical analysis show that the effect of the RG flow is for $P(x, \Omega)$ to generically develop a singularity at small arguments. This is consistent with our intuition that each RG step can only populate the support of the new distribution up to $\Omega/2$, which is a strict upper bound for J' as seen from Eq. (2.43). The universal attractor is approximately a power law,

$$P(x, \Omega) \approx \frac{\alpha}{\Omega} x^{-1+\alpha} \quad (2.45)$$

with a small $\alpha > 0$. This can be verified to be the case if $P_0(J)$ is itself a power law, and is numerically found to hold for other initial distributions as well. Notably, $P_0(J)$ need not be singular for a singular pdf to develop.

So far, the SDRG was applied to the study of the ground energy, or in other words, the zero temperature scenario. A finite- β generalization is possible by applying the same perturbative idea to the free energy $F = \log Z(T)$ instead of the Hamiltonian. The decimation step is then slightly modified, though the spirit of the procedure does not change. It is even possible to include the effects of a nonzero magnetic field, though the resulting equations are much more complicated and elude any analytical considerations. In this way, Dasgupta and Ma were able to show through semianalytical and numerical arguments that their theory predicts a quasi-power-law behavior for the specific heat, the magnetization and the magnetic susceptibility:

$$C = \frac{1}{T} \frac{d^2 F}{dT^2} \sim T^{\gamma_c}, \quad (2.46)$$

$$M^z = \langle S^z \rangle \sim T^{\gamma_M}, \quad (2.47)$$

$$\chi = \frac{1}{L} \frac{dM^z}{dh^z} = \frac{\langle S^2 \rangle}{3LT} \sim T^{\gamma_s-1} \quad (2.48)$$

with γ_c, γ_s and γ_M only containing a weak logarithmic dependence on T .

In Section 2.4.2 we will describe a variant of Dasgupta and Ma's SDRG technique used to study high-energy eigenstates in the same model.

2.3.3 SU(2) symmetry

The random-bond Heisenberg model, Eq. (2.30), possesses an internal SU(2) symmetry group generated by the spin algebra $\{S^\alpha\}_{\alpha=x,y,z}$, where $S^\alpha = \sum_i s_i^\alpha$ is the total spin operator. We now explain in some detail what this entails for the eigenstates of the system. In the next section, we present two arguments suggesting that one of the consequences — namely the presence of *symmetry multiplets* — prevents the coexistence of such a symmetry with

MBL. The arguments naturally extend to $SU(N)$ for any $N \geq 2$ and in fact to all non-Abelian groups.

The $SU(2)$ symmetry of the Hamiltonian manifests itself as the two independent conservation laws²³

$$[H, S^z] = 0, \quad [H, \mathbf{S}^2] = 0, \quad (2.49)$$

which means that we can choose a basis in the Hilbert space of simultaneous eigenstates of H, S^z and \mathbf{S}^2 . The corresponding quantum numbers are the energy E , the *magnetization* M and the *spin* S .

The fixed- (S, M) subspace is in general multidimensional, corresponding to several, linearly independent ways of arranging spins- $\frac{1}{2}$ into a spin- S multiplet. For instance, four spins can be combined in the two following orthonormal ways:

$$|\psi_1\rangle = \frac{|\uparrow\downarrow\rangle - |\downarrow\uparrow\rangle}{\sqrt{2}} \otimes \frac{|\uparrow\downarrow\rangle - |\downarrow\uparrow\rangle}{\sqrt{2}}, \quad (2.50)$$

$$|\psi_2\rangle = \frac{1}{\sqrt{3}} \left[|\uparrow\uparrow\downarrow\downarrow\rangle - \frac{|\uparrow\downarrow\rangle + |\downarrow\uparrow\rangle}{\sqrt{2}} \otimes \frac{|\uparrow\downarrow\rangle + |\downarrow\uparrow\rangle}{\sqrt{2}} + |\downarrow\downarrow\uparrow\uparrow\rangle \right], \quad (2.51)$$

both of which have $S = M = 0$. The source of this degeneracy can be appreciated in full generality by recalling the composition rule for $\mathfrak{su}(2)$ representations [70],

$$\Gamma(s_1) \otimes \Gamma(s_2) = \Gamma(|s_1 - s_2|) \oplus \Gamma(|s_1 - s_2| + 1) \oplus \cdots \oplus \Gamma(s_1 + s_2), \quad (2.52)$$

where $\Gamma(s)$ denotes the unique (up to isomorphisms) representation of $SU(2)$ of spin s , namely, whose dimension is $d = 2s + 1$. Every direct summand denotes a different block in the block-diagonal form of the product representation at the left hand side, *i.e.* we are decomposing the representation in orthogonal invariant subspaces which cannot be further reduced.

Eq. (2.52) describes the *fusion* of two spins s_1 and s_2 . When the total spin s_{12} is specified among all its possible values $|s_1 - s_2|, \dots, s_1 + s_2$, the result of the fusion is a *spin multiplet*, *i.e.* a $(2s_{12} + 1)$ -dimensional vector space on which the quadratic Casimir \mathbf{S}^2 acts as the multiplication by $s_{12}(s_{12} + 1)$. If the total magnetization m_{12} is also specified, a single state is picked out, which may be explicitly realized as the appropriate linear combination

$$|s_{12}, m_{12}; s_1, s_2\rangle = \sum_{m_1, m_2} \text{CG}[s_1 m_1 \otimes s_2 m_2 \rightarrow s_{12} m_{12}] |s_1 m_1\rangle |s_2 m_2\rangle, \quad (2.53)$$

where $\text{CG}[\cdot]$ denotes the Clebsch–Gordan coefficients, which are known exactly for $\mathfrak{su}(2)$. We also employ the *fusion tree* notation

²³We follow the standard convention of conserving the z - rather than the x - or y -component.





Figure 2.4: Two possible fusion tree structures for four particles. Labeling either of these trees in compliance with the $\mathfrak{su}(2)$ fusion rules (2.52) produces orthonormal states via (2.53), and there is a unitary change of basis switching between the two “templates”.



to depict the state (2.53).

A many-particle multiplet can be obtained by combining several fusions, with an L -particle multiplet requiring $L - 1$ fusions. This can be accomplished in several different ways, corresponding to the different ways of decomposing a product of L elementary representations in binary products. Every decomposition gives rise to a distinct geometry for the fusion tree. For example, a four-spin state can be obtained via the fusion

$$\Gamma^{(1/2)} \otimes \left(\left(\Gamma^{(1/2)} \otimes \Gamma^{(1/2)} \right) \otimes \Gamma^{(1/2)} \right) \quad (2.54)$$

(Fig. 2.4a) or

$$\left(\Gamma^{(1/2)} \otimes \Gamma^{(1/2)} \right) \otimes \left(\Gamma^{(1/2)} \otimes \Gamma^{(1/2)} \right) \quad (2.55)$$

(Fig. 2.4b) or in a number of other ways.

Since the tensor product is associative, all fusions must provide the same product representation

$$\Gamma^{(1/2)} \otimes \Gamma^{(1/2)} \otimes \Gamma^{(1/2)} \otimes \Gamma^{(1/2)} = \Gamma(0) \oplus \Gamma(0) \oplus \Gamma(1) \oplus \Gamma(1) \oplus \Gamma(1) \oplus \Gamma(3) \quad (2.56)$$

up to isomorphism, *i.e.* all multiplets encoded by trees with the same total spin and the same number of leaves are unitarily equivalent.²⁴

When the value of the total spin S as well as the leaves are fixed, there is only a specific number of ways one can assign spin values to the internal nodes (called *partial* or *block*

²⁴Leaves in a fusion tree correspond to physical particles, which we always assume to be spins- $\frac{1}{2}$.




Figure 2.5: Basis state representation of a system of 8 spins-1/2 using (a) the tree basis or (b) the computational basis. In (a), the quantum numbers $(s, m; s_1, \dots, s_6)$ describe the spin values of the partial fusions plus the total spin and total magnetization of the system. In (b), the quantum numbers $(\mu_1, \mu_2, \dots, \mu_8)$ describe the individual magnetizations of the physical spins.

spins) consistently with the $\mathfrak{su}(2)$ fusion rules (2.52). Such a labeling, plus the specification of the total magnetization M , singles out a specific state through formula (2.53), and the set of all such states forms an orthonormal basis of the (S, M) sector of the L -particle Hilbert space. This implies that the number of labelings is precisely the dimension of that sector, and in particular it cannot depend on the tree structure. In fact, the number can be computed exactly (Appendix B):

$$D_{L,S,M} = \frac{2(2S+1)}{L+2S+2} \binom{L}{L/2+S}. \quad (2.57)$$

The lack of an M -dependence in $D_{L,S,M}$ should be obvious from the fact that all magnetization subsectors in an S -multiplet have the same dimension (namely, one). In particular, if one is only interested in fixing the total spin but not the magnetization, the resulting dimension is simply $D_{L,S} = (2S+1)D_{L,S,M}$. One can verify that $\sum_S D_{L,S} = 2^L$ (the sum runs from $S = 0$ to $S = L/2$ if L is even, and from $S = 1/2$ to $S = (L-1)/2$ if it is odd).

We call a basis built by fusing spins with a fixed (S, M) a *tree basis*. This is to be contrasted to the computational basis parametrized by the L magnetizations of the physical spins. Both bases are parametrized by L numbers, but unlike the computational basis, the quantum numbers of the tree basis correspond to nonlocal observables — the partial spins plus the total spin and magnetization. See Fig. 2.5 for a pictorial representation.

The advantage of using a tree basis is that all states built off of it will be invariant under the same $SU(2)$ group as the Hamiltonian, which can be exploited in several ways. We mention but two applications.

The first one is numerical: by building an $SU(2)$ -invariant Hamiltonian such as Eq. (2.61) directly in the tree basis, it is possible to restrict oneself to a given (S, M) - or S -sector without

the need to first instantiate the full-space matrix, cutting down the memory requirements by a large factor. For instance, fixing (S, M) to some $O(L^0)$ value causes the space dimension to scale as $D_{L,S,M} \sim 2^L/L^{3/2}$ for large L , to be compared with the scaling $D_{L,M} \sim 2^L/\sqrt{L}$ achieved by only fixing the magnetization (which is a very simple and standard restriction). An important caveat is that it is not always easy to compute brackets of operators between tree states. Indeed, since tree bases are nonlocal, all operators possessing a natural tensor product structure in the computational basis, *e.g.* the spin operators, will not be expressible by means of Kronecker products of smaller operators. Luckily, symmetry considerations plus some heavy-duty algebra make it possible to compute brackets of 1- and 2-spin operators on the tree basis via an efficient recursive formula (see Suppl. Mat. of [212]).

The second one is the observation that the Hamiltonian's eigenstates are expected to be well approximated by tree states, at least in the strong-disorder regime. We will explain exactly what we mean by that in Section 2.4.2, when discussing a generalization of SDRG to the high-energy region of the spectrum. Under this assumption, eigenstates must have a much smaller participation ratio with respect to some tree basis than the computational basis. This will enable us to upper bound the entanglement entropy of the eigenstates much more effectively.

2.3.4 Symmetry and localization

We now discuss whether a $SU(2)$ -symmetric Hamiltonian may display conventional MBL phases, or more accurately, whether the LIOM framework presented in Section 2.2 can coexist with $SU(2)$ multiplets. The analogous question for Abelian symmetries was answered in the positive when it was established that a \mathbb{Z}_2 -symmetric Hamiltonian such as the disordered transverse-field Ising model [126], as well as some non-free-fermion augmentations of it through next-to-nearest neighbor interactions [204] or a nearest-neighbor transverse perturbation [145], can localize while preserving the spin-flip symmetry. At the same time, all these models were shown to exhibit spin-glass ordering at low temperature and strong disorder: as the relative strength of the spin-spin interaction is increased, the \mathbb{Z}_2 symmetry is broken in the eigenstates without spoiling localization — an example of MBL-to-MBL dynamical transition.

A different picture emerges when considering non-Abelian symmetries such as $SU(2)$. It was argued by Potter and Vasseur that it is impossible for a quantum state to display *any* kind of non-Abelian symmetry while also admitting LIOMs [209]. The reason for this is that non-Abelian symmetries come with nontrivial irreducible representations, namely multiplets. Assume that LIOMs exist alongside a symmetry group G , acting on the Hilbert space $\mathcal{V}^{\otimes L}$ via the appropriate product representation (here \mathcal{V} is assumed to be a nontrivial single-particle representation space, *e.g.* the fundamental representation of $\mathfrak{su}(2)$ for spin- $\frac{1}{2}$ particles). Compatibility means that LIOMs and group generators commute with each other; this implies that it is possible to find a system of eigenstates which are

simultaneously labeled by a given set of LIOMs $\{I_i\}_{i=1}^N$ and form a multiplet under the action of the symmetry group. Alternatively, symmetry multiplets can be labeled by the *constant* values taken by the LIOMs on them:

$$|\psi\rangle = |I_1, \dots, I_N; f\rangle \in \mathcal{V}_{I_1, \dots, I_N}, \quad (2.58)$$

where $\mathcal{V}_{I_1, \dots, I_N}$ is a representation space of dimension D and f spans the multiplet's degeneracy, $f \in \{1, \dots, D\}$.

Given such a state $|\psi\rangle$, and provided $D > 1$, it is possible to act on it via a symmetry transformation which will transform it into a distinct state (*i.e.* act on the quantum number f) without changing its energy or any of the LIOMs. Moreover, the symmetry action is *local*. This can be proven by the following argument: suppose there exists a reference eigenstate $|\psi_0\rangle$ which is symmetry-invariant, *i.e.* living in a $D = 1$ singlet. Then, changing the value of a single LIOM, say $I_a \mapsto I'_a$, will in general define a $D > 1$ multiplet, whose states will only differ from $|\psi_0\rangle$ on the support of the a -th LIOM, which is exponentially localized by definition. Therefore, whatever the symmetry does on the states of this multiplet, it can only change them within that support, *i.e.* the symmetry acts *locally* on the eigenstates.

The above argument entails that in the simultaneous presence of multiplets and LIOMs, local degeneracies are bound to emerge. In particular, as the number of LIOMs is extensive, we expect an exponential number of them. This generically prevents a putative localized state from being stable. The conclusion is that non-Abelian symmetries must either break or be broken by localization. Notice that Abelian symmetries are safe from this argument as then $d = 1$ always.

An alternative argument proceeds from considering the bipartite entanglement entropy of multiplet states [212]. Consider two separate spin chains which are joined together by switching on a local coupling at $t = 0$. Assuming $SU(2)$ symmetry, each eigenstate of the disconnected system is a product of two states belonging to two separate multiplets,

$$|\psi(t < 0)\rangle = |\psi_L\rangle \otimes |\psi_R\rangle, \quad |\psi_{L/R}\rangle \in \mathcal{V}_{L/R}, \quad (2.59)$$

of respective spins s_L and s_R . With no loss of generality, we can suppose that $s_L \geq s_R$. After switching on the coupling (which is assumed to respect the global symmetry), the new eigenstates must transform as a new multiplet with some definite value S of the total spin, $s_L - s_R \leq S \leq s_L + s_R$, and magnetization M . We can think of $|\psi(t < 0)\rangle$ as a pair of fusion trees which are later conjoined at the top level into a single tree at $t = 0$.²⁵

The generic fusion of two multiplets is a state with $O(s_R)$ nonzero coefficients, so one should expect an $O(\log s_R)$ value for the half-chain entanglement entropy. Since a typical

²⁵In general, a fixed- (S, M) eigenstate will be obtained by joining several different combinations of left- and right-multiplets. However, considering more than one left-right pair can only increase the entanglement entropy. As we are interested in lower-bounding the entropy, we can restrict ourselves to the simplest case of a single multiplet pair.

eigenstate has $s_R \in O(L^{1/2})$, this implies a generic $O(\log L)$ scaling of the entanglement entropy, violating the area law predicted by the LIOM scenario. Notice that, in the assumption that fusion trees are true eigenstates of the system, the partially-fused spins (*i.e.* the spins labeling the internal nodes of the trees) are an extensive set of exact constants of the motion, playing a role similar to LIOMs in actual MBL systems. The difference is that, while many partial spins have a small spatial support — namely, the ones located close to the leaves — we need some highly nonlocal spins as well in order to form the whole tree. For instance, the top spin (the one “trivially conserved” by the global symmetry) has a support spanning the entire chain. In fact, we show in Appendix C.2 that the *average* support size of a partial spin in a random fusion tree is $O(\log L)$.

It is worth pointing out that the outlined arguments are both referring to situations where the eigenstates under considerations retain the full symmetry of the system (or at least a non-Abelian subgroup thereof). In general, the presence of a non-Abelian symmetry group G for the Hamiltonian only means that it is *possible* — not mandatory — to choose a G -symmetric eigenbasis. However, less symmetric eigenstates may be constructed as well, to which the above arguments do not apply.

For example, consider a pair of uncoupled Anderson chains, with identical disorder on the corresponding sites:

$$H = -t \sum_{i,s} \left(c_{i,s}^\dagger c_{i+1,s} + c_{i+1,s}^\dagger c_{i,s} \right) + \sum_{i,s} \epsilon_i c_{i,s}^\dagger c_{i,s}, \quad (2.60)$$

where s is a label for the two chains, say $s \in \{\uparrow, \downarrow\}$. The model Eq. (2.60) has an $SU(2)$ “flavor” symmetry that mixes the two kinds of fermions, just as seen for the Hubbard model Eq. (2.32) (from which Eq. (2.60) is obtained by setting $U = 0$). However, it is clearly possible to construct eigenstates by just taking the product of two eigenstates of the two different chains, both of which are known to be Anderson localized for any value of the disorder. In this way, “trivially localized” eigenstates of the model can be exhibited in spite of the $SU(2)$ invariance.

A less trivial case of localized states in $SU(2)$ symmetric models can be found in the *interacting* Hubbard model, which in fact admits *exponentially large* invariant subspaces within which the dynamics is Anderson localized [127]. In both cases, no contradiction with our stated arguments obtains, as the eigenstates undergoing localization are always contained in a sector of the Hilbert space where the non-Abelian symmetry, due to the presence of a global \mathbb{Z}_2 ($c_{i,\uparrow} \leftrightarrow c_{i,\downarrow}$) which is incompatible with the full rotational symmetry, is broken down to an Abelian $U(1)$.

From the two arguments presented above, not to mention experimental [224] and numerical [210, 268] evidence for the existence of non-ergodic behavior in the disordered Hubbard model, Eq. (2.32), we conclude that highly symmetric systems must fail to fully localize. That said, it is not clear what kind of phenomenology we should expect as a consequence of this failure. Do these systems eventually thermalize, or is it possible

that some exotic phase emerges, characterized by strongly subthermal (yet non-area-law) eigenstate entropy and slow dynamics? In order to shed some light on the nature of systems with symmetry-enforced degeneracy and seek at least a partial answer to the above questions, we now delve into a thorough investigation of the random-bond Heisenberg model.

2.4 Case study: the random-bond Heisenberg model

We will see that all our findings point toward the existence in this model of a broad, peculiar non-ergodic regime in which the system's eigenstates are well described by eigenstates which obey neither ETH nor the usual LIOM picture. The system is expected to eventually thermalize through the emergence of multispin resonances, but at such length and time scales as to be practically inaccessible by cold-atom experiments. The proposed picture for the regime, including the crossover to ETH, is illustrated in Fig. 2.6.

2.4.1 Preliminary definitions and remarks

The model we study is the one-dimensional random-bond Heisenberg model (RBHM), described by the Hamiltonian

$$H = \sum_i J_i \mathbf{s}_i \cdot \mathbf{s}_{i+1}. \quad (2.61)$$

The couplings J_i are quenched random variables identically and independently distributed on $[-1, 1]$ according to the probability distribution

$$P(J) = \frac{\alpha \theta_H(1 - |J|)}{2|J|^{1-\alpha}}, \quad (2.62)$$

(with θ_H the Heaviside theta function) which we average over using double brackets:

$$\llbracket X \rrbracket = \int \left(\prod_i P(J_i) dJ_i \right) X(J_1, \dots, J_L). \quad (2.63)$$

As discussed in Section 2.3.2 (cf. in particular Eq. (2.45)), a distribution of the form Eq. (2.62) emerges in a natural way when considering the effective low-temperature regime. A noteworthy difference between our choice of distribution and the Dasgupta–Ma universal attractor is that our pdf is symmetric around $J = 0$ rather than purely antiferromagnetic; in more generality, we could use a parameter $0 \leq \eta \leq 1$ to tune the fraction of AFM-to-FM couplings in the system. However, since we are interested in the properties of highly excited eigenstates in the middle of the many-body band, where entropy completely dominates energy, we do not expect these properties to exhibit a significant dependence on the choice of η ,²⁶ so our choice falls on the unbiased value $\eta = 0.5$.

²⁶It is known that not even the ground state properties depend strongly on the value of η , except for the

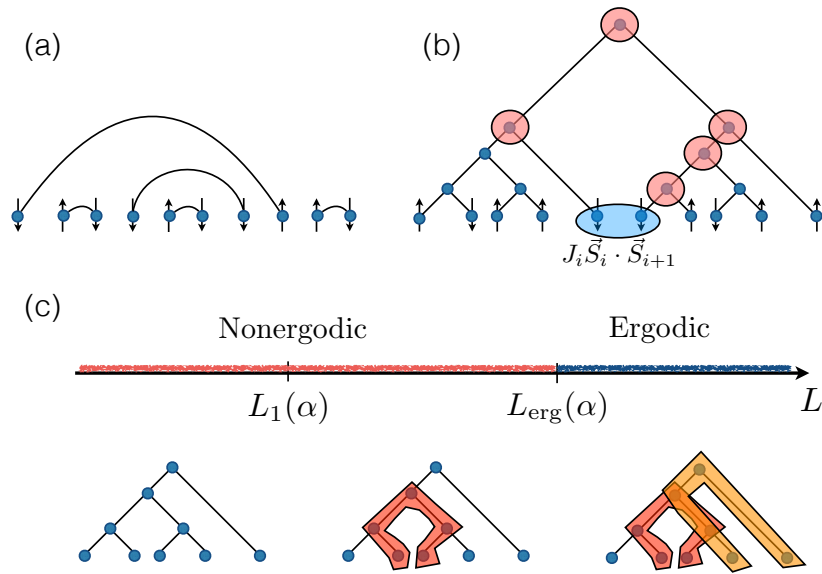


Figure 2.6: (a) A cartoon of the ground state of a random antiferromagnetic Heisenberg chain; (b) Strong-disorder renormalization group aims to construct approximate eigenstates. It yields a tree state, characterized by its geometry and the choice of total block spins at each node (see main text). The Heisenberg Hamiltonian, written in this basis, gives rise to processes which can change the spins along the *fusion path* connecting two neighboring spins, to one of the block spins; (c) A schematic dynamical phase diagram of the random Heisenberg model. There are three regimes: (I) at short length scales, $L < L_1(\alpha)$, the SDRG tree states are accurate approximations to the eigenstates; (II) at intermediate length scales, $L_1(\alpha) < L < L_{\text{erg}}(\alpha)$, there are resonances but the system remains non-ergodic; (III) above some large length scale $L > L_{\text{erg}}(\alpha)$, the resonances proliferate and the system becomes thermalizing (see Section 2.4.4 for a definition of these scales).

The parameter $\alpha > 0$, which determines the power-law tail behavior of $P(|J|)$, also controls the strength of the quenched disorder in the following sense: one can see from Eq. (2.62) that the ratio of two neighboring couplings in the system has a typical value

$$\frac{|J_{1\vee 2}|}{|J_{1\wedge 2}|} \Big|_{\text{med}} = 2^{1/\alpha}, \quad (2.64)$$

where $J_{1\vee 2}$ ($J_{1\wedge 2}$) is the maximum (minimum) between J_1 and J_2 and X_{med} is the median of X .

The ratio (2.64) increases exponentially when $\alpha \rightarrow 0^+$. Therefore, at small α it becomes more and more likely to find exchange constants in the system that are *much* larger than the two neighboring ones. This is exactly the “strong disorder” condition that enables SDRG, as we discuss below.

Another quantity of interest is the smallest coupling J (in absolute value) in the whole system, representing the “weakest link”. We find that

$$\min_i |J_i| \sim \frac{\Gamma(1/\alpha)}{\alpha L^{1/\alpha}}. \quad (2.65)$$

For $\alpha = 0.3$ and $L \approx 20$, this coupling can be as small as $10^{-3} \llbracket J \rrbracket$.

Throughout the paper it will be helpful to contrast our findings to the properties of the random field XXZ model, which has been studied extensively in the literature (see Refs. [2, 14, 199] for recent reviews):

$$H_{\text{XXZ}} = t \sum_i (s_i^x s_{i+1}^x + s_i^y s_{i+1}^y) + U \sum_i s_i^z s_{i+1}^z + \sum_i h_i s_i^z \quad (2.66)$$

The model (2.66) can be mapped, via Jordan–Wigner transformation (Appendix A), onto an interacting fermionic problem with t representing the hopping amplitude, U the nearest-neighbor interaction, and h_i the random on-site potential with a variance that we denote by W . In the following, for concreteness, we will assume that $t \sim U$, such that the disorder strength is described by a single dimensionless parameter $L_W = W/t$. Then, the XXZ model is known to have a diffusive–subdiffusive dynamical transition at $L_W \approx 0.55$ [270] and an MBL–thermal transition at $L_W \approx 3.5$ [162].

In the limit of strong disorder, $W \gg t$, the parameter L_W can be interpreted as a typical distance between (rare) pairs of “resonant” sites in the model that happen to have close enough values of the magnetic field to enable resonant spin exchange (or, equivalently, hopping in the fermionic model). A resonance between spins 1 and 2 appears, for example, if $|h_1 - h_2| \lesssim t$. Starting at a very large disorder these resonant sites are typically well separated by distances of $O(L_W)$, and one can show that they will not mix at any order of perturbation theory [128, 129]. By “mixing” we mean that the resonant pairs can exchange

extremal points $\eta \in \{0, 1\}$ [263].

energy and become strongly entangled in the eigenstates. The fact that resonances are rare and isolated at $L_W \geq 3.5$ is intimately related to the low, area-law entanglement scaling of eigenstates, and the existence of a complete set of LIOMs [125, 219, 227, 228, 260]. As the disorder strength is decreased, the resonant pairs of spins eventually become mixed, forming a connected network; then, LIOMs are destroyed, becoming nonlocal, and the system exits the MBL phase.

What is the proper *quantitative* measure of disorder strength in an SU(2) symmetric spin chain? The estimate (2.64) for the *typical* ratio of the neighboring couplings in a Heisenberg chain suggests that the disorder experienced by the system becomes exponentially large in $1/\alpha$. Therefore, naively one could expect that, similar to the case of the random field XXZ chain, where $L_W \propto W$, a length scale $L \propto e^{1/\alpha}$ (the inverse of the typical ratio of neighboring couplings) would determine the density of rare resonances. However, as we show in Section 2.4.4, another measure of disorder is in fact important. Specifically, one can introduce a length scale $L_1(\alpha)$ with a meaning similar to that of L_W in the XXZ model, Eq. (2.66): $L_1(\alpha)$ defines a typical distance between local resonances in the system. This length scale diverges when α goes to zero but, in contrast to the typical ratio of couplings, Eq. (2.64), only *polynomially* in $1/\alpha$. Our numerical findings below are consistent with $L_1(\alpha) \propto \alpha^{-0.4}$.

If the usual MBL scenario applied here, some $L_c \in O(L^0)$ would exist such that, if $L_1(\alpha) \geq L_c$ the resonances would not proliferate and the novel non-ergodic phase would be stable. Instead, the entanglement pattern of eigenstates, as well as the nonlocal nature of some integrals of motion induced by SU(2), lead to the eventual proliferation of resonances at *any* disorder strength, and so for *any* value of $L_1(\alpha)$, provided the system is sufficiently large. Thus, another scale marking the crossover from the localized to the ergodic phase emerges. We denote this length scale, where ergodicity is restored, by $L_{\text{erg}}(\alpha)$. Below, we provide strong evidence for the delocalization scenario described above. In systems with relatively weak disorder, the length scale $L_{\text{erg}}(\alpha)$ manifests itself *e.g.* in the level statistics and ETH violation for matrix elements of local observables that we study via exact diagonalization (see Section 2.4.3 for details). At stronger disorder, no tendency towards ergodicity restoration can be observed in ED studies due to size limitations. However, a detailed analysis of the resonant processes (see Section 2.4.4) allows us to estimate $L_{\text{erg}}(\alpha)$ in this case as well.

2.4.2 SDRG for excited states

In this section, we qualitatively describe the SDRG approach to disordered Heisenberg chains — redirecting the interested reader to the relevant literature — and discuss the properties of the tree states it yields. We emphasize that such states differ from conventional MBL states in two crucial aspects: first, they have a parametrically larger entanglement entropy, and second, one cannot define a complete set of LIOMs for them.

A very large typical ratio of two neighboring couplings found for small α , Eq. (2.64),




Figure 2.7: A multiplet of eigenstates predicted by the SDRG for a system of 12 spins- $\frac{1}{2}$. The leaves of the tree represent the physical spins. The tree describes the way these spins are fused into larger block spins over the course of the SDRG. The numbers in the nodes indicate the resulting partial spins. The value in the top node (marked red) is the total spin S of the system ($S = 1$ in the present example). S is an exact integral of motion. $(2S + 1)$ different states in the multiplet can be distinguished by additionally specifying the projection of the spin in the top node to the z -axis.

suggests that the properties of the system can be described using the SDRG framework. The idea of SDRG is to identify a local “grain” in the system that is strongly coupled inside, but, due to strong disorder, only weakly coupled to the rest of the system. The state of the grain is then approximated by one of the eigenstates of its Hamiltonian, with the rest of the system decoupled. If one is looking for the ground state, the eigenstate of the grain is chosen to be its ground state. Alternatively, if one is interested in constructing a random highly excited eigenstate that is effectively at infinite temperature — as we are in the present discussion — some eigenstate of the grain is randomly chosen. Then, the effective Hamiltonian of the system in which the grain is in the chosen eigenstate (or, more generally, a multiplet of states if symmetries dictate degeneracies in the spectrum of the grain’s Hamiltonian) is calculated by perturbation theory in the grain–system coupling.

One can then iterate this procedure, assuming that the disorder in the effective Hamiltonian remains strong. This is indeed the case for, *e.g.*, ground states of random antiferromagnetic (AFM) Heisenberg chains [106]. Then, a repeated application of the SDRG rules results in an approximate wave function of the whole system, obtained by “patching together” the wave functions of the grains.

An in-depth discussion of the SDRG rules for excited states of the Heisenberg chain can be found in Refs. [8, 213], to which we refer for the details. Qualitatively, for this system a grain is a pair of neighboring spins coupled by a strong bond; its eigenstates (which come in $SU(2)$ multiplets) are labeled by the total spin of the grain. The SDRG procedure replaces such spin pairs by effective (typically larger) spins, *i.e.* it assigns some total spin to larger and larger blocks of spatially contiguous spins in the system. The resulting approximation

for an eigenstate (more precisely, for a degenerate symmetry-enforced multiplet²⁷) is a kind of a tree tensor network, illustrated in Fig. 2.7. The nodes of the tree represent the block spins identified in the SDRG process. The structure of the tree reflects the order in which the elementary spins of the system should be fused to give an (approximate) eigenstate.

The fusion of spins in the course of the SDRG must be supplemented by a perturbative account of the interaction of merging spins with the rest of the system. In the present setting of an infinite-temperature SDRG, where spins typically fuse into non-singlet states, a first-order perturbation theory (that simply amounts to the projection of the fusing spins onto the direction of the total spin) suffices in most cases. The resulting renormalization of couplings is weaker than the one that occurs in the low-temperature SDRG for AFM spin chains, where the spins always fuse into singlets, and therefore a second-order perturbative treatment is required to find new renormalized couplings (see Appendix A in [213]). Still, the distribution of couplings developed in the course of SDRG turns out to be broad (see Ref. [8] and below).

As mentioned in Section 2.3.4, within the SDRG approximation the values of the block spins label the eigenstates of the Hamiltonian and bear similarity to the LIOMs of the conventional MBL phase. An eigenstate of the Hamiltonian is also an eigenstate of a sequence of these operators, just as an eigenstate of an MBL Hamiltonian is simultaneously an eigenstate of each LIOM. However, there are two major differences between these quantum numbers and LIOMs.

First, in the MBL phase the eigenstates of H are at the same time eigenstates of a *fixed set* of LIOMs. Total spins of the blocks in our problem would form conserved operators if different eigenstates were represented by *geometrically identical* trees, which differ only in the values of the block spins. In reality, the order in which spins are merged over the course of SDRG depends not only on the particular disorder realization, but also on the eigenstate of the grain, which is randomly picked at any given step of the SDRG. Thus, the values of the block spins, in general, cannot be promoted from labels of a particular eigenstate to operators acting in the full Hilbert space. The structure of larger blocks depends on the history of choosing total spins at the earlier steps of SDRG.

Second, LIOMs in an MBL system are quasilocal, exponentially localized in space operators [126,227,228]. In contrast, the block spins of the strongly disordered Heisenberg chain have a hierarchical structure. While some of them (living near the bottom of the tree) can be expressed in terms of an $O(L^0)$ number of the original spin operators \mathbf{s}_i , the other ones, found at the higher levels of the tree, are highly nonlocal in terms of the original spins. Thus, $SU(2)$ symmetry forces some integrals of motion to become nonlocal. Therefore, SDRG (in the regime of its validity) describes a non-ergodic phase of a new kind, with a *partial*, rather than complete set of LIOMs. Our goal is to investigate the

²⁷As most properties of the states comprising a multiplet are actually independent of the particular state and depend on the multiplet alone, throughout our discussion we often refer to SDRG trees as specifying a single quantum state. One may assume, for example, that in each multiplet we focus on the state with $M = 0$ ($1/2$) if the length of the system is even (odd).

stability of this putative phase.

The peculiar non-ergodic character of tree eigenstates manifests itself in the scaling of entanglement entropy. For simplicity, we will consider the entanglement entropy of an eigenstate with respect to the a cut in the middle of the chain,

$$\mathcal{S}_{\text{ent}}(L/2) = -\text{Tr}(\rho_{L/2} \log_2 \rho_{L/2}), \quad (2.67)$$

where $\rho_{L/2}$ is the reduced density matrix of half-chain in the chosen eigenstate and the trace is taken over the degrees of freedom in the other half of the system. A bound for the entanglement entropy depends on the tree structure describing a given state, in particular, on the tree height h (the number of levels between the very top node of the tree and the physical spins). We find, via numerical simulations, that typical states produced by the SDRG procedure have a logarithmic height, $h \propto \log L$. It is then possible to show (see Appendix B in Ref. [213]) that the entanglement entropy of a single typical²⁸ tree satisfies

$$c_1 \log_2 L \lesssim \mathcal{S}_{\text{ent}}(L/2) < c_2 \log_2^2 L, \quad (2.68)$$

where c_1 and c_2 are numerical constants of order unity that depend on the statistical properties of the tree. Thus, the entanglement of tree states scales faster than the area law found in MBL, but significantly slower compared to the thermal entanglement for an infinite-temperature state, $\mathcal{S}_{\text{th}}(L/2) \approx L/2$ (measured in bits).

The upper bound on the entanglement entropy in Eq. (2.68) can also be generalized to the case when the state in question is not a single tree state but rather a linear combination of n_T tree states:

$$\mathcal{S}_{\text{ent}}(L/2) < c_2 \log_2^2 L + \log_2 n_T. \quad (2.69)$$

Although this bound might seem weak, it has an important implication, which will be used below: if the system's eigenstates become ergodic, they must be represented by an exponentially large number of tree states.

Validity of SDRG and (in)stability of the tree states

The SDRG is a heuristic real-space renormalization procedure relying on strong disorder. The tree states generated by SDRG are not exact eigenstates of the Heisenberg spin chain, but how accurate are they? Historically, at each step of SDRG one checks that the disorder in the effective Hamiltonian remains strong, such that strong couplings can be found; one can then check for the absence of resonances involving a small number of spins, to make sure that the neglected processes do not destroy the tree structure. While for the analysis of ground states this is often sufficient, it is unclear whether such tests can guarantee the accuracy of SDRG for the excited states.

²⁸The upper bound on the entanglement entropy in Eq. (2.68) holds in fact for all the states described by trees of logarithmic depth.

Below we will check the validity of SDRG for excited states using several approaches. First, we will compare SDRG tree states to the exact eigenstates for system sizes up to $L = 26$, obtained numerically. We will use a number of measures, such as level statistics, and the eigenstate thermalization hypothesis (ETH) and its breakdown. Second, to describe large system sizes, we will develop an approach to account for many-body processes that are usually neglected in SDRG, and to test their relevance. We introduce this approach qualitatively now, and we will apply it in what follows.

Suppose that SDRG yielded some tree state $|\Psi_{\text{RG}}^0\rangle$, specified by the tree geometry and the choice of partial spins in each node. Instead of considering the effective Hamiltonian at every step, we can write the original Hamiltonian exactly in the basis of tree states with the geometry identical to that of $|\Psi_{\text{RG}}^0\rangle$. The first key observation is that the selection rules imposed by symmetry facilitate the analysis of the relevant processes; more specifically, the block spins along the fusion path of two contiguous physical spins can only change by $\Delta S = 0, \pm 1$.²⁹ The second observation is that, given that the typical spins of larger blocks grow (as the square root of the block size³⁰), the tree states connected to $|\Psi_{\text{RG}}^0\rangle$ by the Hamiltonian are expected to have the same geometrical structure. This is because for large spins, strong bonds remain strong when a value of some spins is changed by $\Delta S \ll S$.

We search for *resonances* between different tree states and characterize their properties. Solving the full eigenvalue problem for large L is hopelessly complicated; thus, we focus on low-order resonances. Effectively, we check whether the Hamiltonian hybridizes a given tree state with its neighbor, say $|\Psi_{\text{RG}}^1\rangle$ (a neighbor is a state such that $\langle \Psi_{\text{RG}}^1 | H | \Psi_{\text{RG}}^0 \rangle \neq 0$). As long as the probability of finding resonances is sufficiently low, we expect the true eigenstates to be localized in the tree basis. This corresponds to a non-ergodic phase (or regime, if it only occurs for sufficiently small system sizes). Alternatively, if there are many resonances which proliferate, it is natural to expect the SDRG to break down and the system to become ergodic.

It is instructive to draw parallels with the conventional MBL phase of the strongly disordered XXZ spin chain in a random magnetic field. The caricature of MBL eigenstates is just product states with a well-defined s_i^z projection for each spin. While corrections to this picture certainly exist (*e.g.* LIOMs are not strictly equal to s_i^z operators) we know that MBL is stable, if the disorder is sufficiently strong. Our aim is to understand whether for SU(2)-symmetric chains tree states, with their built-in correlations and unusual entanglement properties, can be stable, representing a dynamical phase distinct from both MBL and ergodic phase.

Below we will use the above approach to reveal a broad non-ergodic regime where tree states are indeed stable. We will also provide evidence that trees eventually become unstable above a certain (α -dependent) system size for all values of α that we study. We

²⁹Additionally, spins should be non-negative, and the $0 \rightarrow 0$ transition is forbidden, as is well-known from optics. These rules all follow from the fusion of spins, Eq. (2.52).

³⁰This is for purely entropic reasons. At infinite temperature, the probability to pick a given spin S is proportional to $D_{L,S}$, which is maximal at $S \in O(\sqrt{L})$.

L	full	$M = 0$	ratio to full	$S = M = 0$	ratio to full
10	1 024	252	24.61%	42	4.10%
12	4 096	924	22.52%	132	3.22%
14	16 384	3 432	20.95%	429	2.62%
16	65 536	12 870	19.64%	1 430	2.18%
18	262 144	48 620	18.55%	4 862	1.85%
20	1 048 576	184 756	17.62%	16 796	1.60%
22	4 194 304	705 432	16.82%	58 786	1.40%
24	16 777 216	2 704 156	16.12%	208 012	1.24%
26	67 108 864	10 400 600	15.50%	742 900	1.11%

Table 2.2: Direct comparison of the Hilbert space dimension of a system of L spins- $\frac{1}{2}$, without restriction and restricted to the $M = 0$ and $(S, M) = (0, 0)$ sectors. The general formulas are respectively given by $D_L = 2^L$, $D_{L,M} = \binom{L}{L/2+M}$ and Eq. (2.57). See also Fig. B.1 and the surrounding discussion for further comparison.

therefore propose the picture that, while for finite systems the dynamics is non-ergodic at strong disorder, ETH should be recovered in the thermodynamic limit (see Fig. 2.6).

2.4.3 Exact diagonalization

We now present our numerical results from exact diagonalization. The code was implemented in collaboration with Rajat K. Panda under the supervision of Antonello Scardicchio at ICTP, and run on a CINECA supercomputing cluster. Large sizes ($L = 26$) were made accessible by a powerful combination of parallelization and the dimensional reduction from constructing the Hamiltonian within the fully SU(2)-symmetric $(S, M) = (0, 0)$ sector of the Hilbert space³¹ (see Table 2.2). However, we omit large- L data when the statistics is insufficient for the analysis. In particular, for $L = 26$ we could only accumulate about ten disorder realizations, therefore that system size is expunged from almost all data.³² All the other system sizes have at least hundreds of samples per data point.

The SRDG and resonance-counting procedures, including the following stability analysis of tree states up to $L = 20$, were implemented and run by Ivan Protopopov under the supervision of Dmitry Abanin.

³¹The $(S, M) = (1, 0)$ sector was also probed, with qualitatively identical results.

³²The memory requirement for mid-spectrum diagonalization at large Hilbert dimensions d is very demanding, scaling approximately as $f(d) \sim ad^b$ with $a = 3.4 \times 10^{-8}$ GiB, $b = 1.9$. The actual memory usage totaled roughly seven terabytes for an $L = 26$ instance.

Probing the stability of tree states

To analyze the accuracy of the SDRG procedure, we first study the participation ratios of exact eigenstates of the system (Eq. (2.61)) in the tree basis generated by SDRG. More precisely, for a given disorder realization $\{J_i\}_i$, we first run the SDRG to generate a single tree state $|\Psi_{\text{RG}}^0\rangle$ with a total spin S . A complete basis of states in the sector with a given total spin S (and total magnetization M) can be built out of $|\Psi_{\text{RG}}^0\rangle$ by fixing the geometry of the underlying tree, but allowing the block spins in the tree (apart from the top one, S) to take all possible values consistent with the angular momentum addition rules. We denote the basis obtained in this manner by

$$\{|\Psi_{\text{RG}}^a\rangle\}_{a=0,\dots,D_{L,S,M}-1} \quad (2.70)$$

where L is the length of the chain and $D_{L,S,M}$ is the dimension of the sector, given by Eq. (2.57).

The state with an index $a = 0$ is the original SDRG state, $|\Psi_{\text{RG}}^0\rangle$. In general, due to the correlations between the geometric structure of the tree and the values of the block spins discussed in Section 2.4.2, many of the states in the basis (with indices $a > 0$) would *not* be approximate eigenstates constructed by the SDRG. We expect, however, that at strong disorder the geometry of the tree that corresponds to the state $|\Psi_{\text{RG}}^0\rangle$ is also appropriate for a number of other SDRG states that do not differ too much from $|\Psi_{\text{RG}}^0\rangle$ in the values of the block spins. In that case, a significant fraction of $|\Psi_{\text{RG}}^{a>0}\rangle$ are in fact “SDRG eigenstates”.

We then perform an exact diagonalization of the Hamiltonian in the basis (2.70).³³ Among all the eigenstates of the Hamiltonian we focus on a single one, denoted by $|E\rangle$, that has maximum overlap with a given $|\Psi_{\text{RG}}^0\rangle$. The quality of $|\Psi_{\text{RG}}^0\rangle$ as an approximation to $|E\rangle$ can be quantified by the inverse participation ratio (IPR) of the state $|E\rangle$ in the SDRG basis (2.70):

$$I_E = \sum_{a=0}^{D_{L,S,M}-1} |\langle E | \Psi_{\text{RG}}^a \rangle|^4, \quad (2.71)$$

and by its inverse $N_E = 1/I_E$, which can be viewed as the number of tree states $|\Psi_{\text{RG}}^a\rangle$ (of a given geometry) that one needs in order to represent the eigenstate $|E\rangle$. Thus, small values of $N_E \sim 1$ indicate that the SDRG is accurate, while a very large $N_E \gg 1$ signals instability of the tree states.

Computing the participation ratio N_E for 10^3 disorder realizations $\{J_i\}_i$ (and a single random SDRG state $|\Psi_{\text{RG}}^0\rangle$ per realization), we investigate the statistical properties of this quantity. We performed numerical simulations for the disorder parameter α ranging from $\alpha = 1.2$ (weak disorder) to $\alpha = 0.3$ (strong disorder). The results are summarized in Figs. 2.8 and 2.9. Figure 2.8 shows examples of the distribution of $\log_{10} N_E$ for different system sizes and two different disorder strengths. We observe that in short systems,

³³Explicit expressions for the matrix elements of the operators $\mathbf{s}_i \cdot \mathbf{s}_j$ in a tree basis were derived in Ref. [212].

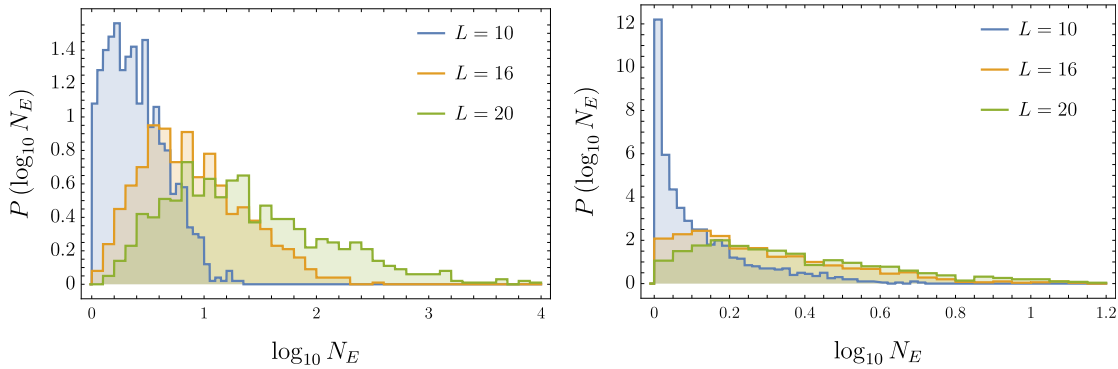


Figure 2.8: Statistics of $\log_{10} N_E$ for $\alpha = 1$ (left panel) and $\alpha = 0.3$ (right panel). Different curves correspond to different system sizes, $L = 10$, $L = 16$ and $L = 20$ (see legend).

$L = 10$, $|\Psi_{\text{RG}}^0\rangle$ is very close to an exact eigenstate even for weak disorder, $\alpha = 1$, in the sense that $N_E \sim 1$. Upon increasing the system size N_E grows, signaling that approximating the eigenstate $|E\rangle$ with a tree state $|\Psi_{\text{RG}}^0\rangle$ becomes less accurate.

The evolution of the typical value of N_E (defined as $e^{\llbracket \log N_E \rrbracket}$) with the system size is illustrated in the top panel of Fig. 2.9. Interestingly, even in the weak disorder regime, $\alpha = 1$, and for the largest system size $L = 20$, the typical $N_E \approx 25$ remains small compared to the dimension of the Hilbert space $D_{L,S,M}$. The latter depends on the spin sector S , which is chosen at random in the present analysis. The SDRG procedure we use generates states with different S in accordance with their probability in the infinite temperature ensemble, $P(S) \propto (2S + 1)D_{L,S,M}$. For $L = 20$ the most frequently encountered value of S is 3, corresponding to the Hilbert space dimension $D_{20,3,0} = 38760$. Moreover, for 90% of the SDRG states $S \leq 5$ and $D_{20,S,M} \geq 10659$. The length dependence of the typical Hilbert space fraction occupied by the energy eigenstate $|E\rangle$ in the tree basis, $e^{\llbracket \log(N_E/D_{L,S,M}) \rrbracket}$, is shown in the bottom panel of Fig. 2.9.

It is instructive to compare the above findings to the behavior of the IPR in the product state basis for the conventional MBL phase. Viewing MBL as a kind of Anderson localization in the Hilbert space, one might naïvely expect that in the MBL regime the eigenstates would exhibit system-size independent IPR, $N_E \gtrsim 1$. It is known [14, 36, 77, 162], however, that in reality MBL eigenstates are rather *fractal* when viewed in the product-state basis: the participation ratio N_E scales as $N_E \propto d^\gamma \propto 2^{\gamma L}$ with an exponent that depends on disorder strength. The fractal behavior stems from perturbative corrections, and resonances discussed at the end of Section 2.4.1 (or, equivalently, it is due to the fact that local integrals of motion have support over more than one lattice site). In the strong disorder limit, $\gamma \propto t/W = 1/L_W \ll 1$. The MBL transition is thus marked *not* by the emergence of the growth of N_E with the system size, but rather by a jump of the exponent γ to its thermodynamic value, $\gamma = 1$ (at infinite temperature).

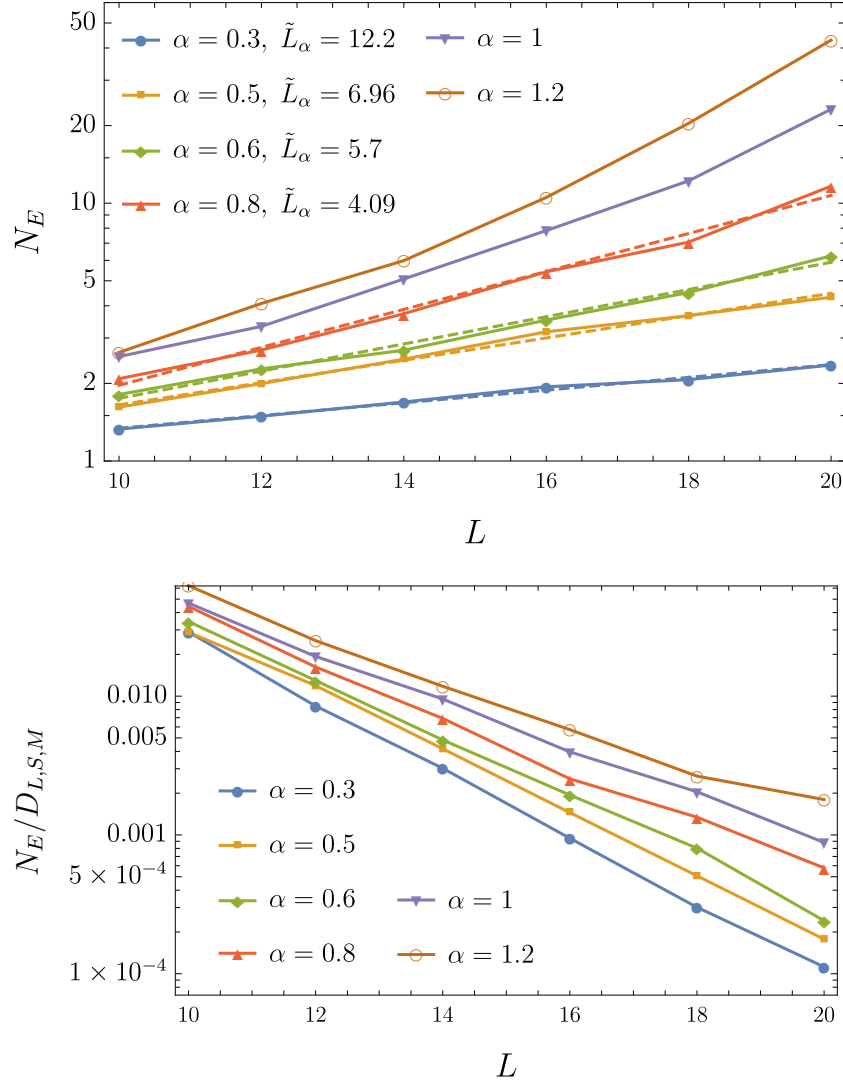


Figure 2.9: Typical number N_E of tree states participating in the eigenstate $|E\rangle$ (top) and the typical value of the fraction $N_E/D_{L,S,M}$ (bottom) versus the system length for different strengths of disorder (see legend). The dashed lines in the top panel represent the exponential fits $N_E \propto 2^{L/\tilde{L}_1(\alpha)}$.

The behavior shown in Fig. 2.9 for the Heisenberg chain is qualitatively similar. At strong disorder, $\alpha \leq 0.8$, the dependence $N_E(L)$ for the available system sizes can be approximated by an exponential fit, $N_E \propto 2^{L/\tilde{L}_1(\alpha)}$ (see dashed lines in the top panel of Fig. 2.9; the corresponding values of the fitting parameter $\tilde{L}_1(\alpha)$ are indicated in the legend). The length $\tilde{L}_1(\alpha)$ grows as the disorder strength is increased. By analogy with the conventional MBL, we can expect the length scale $\tilde{L}_1(\alpha)$ to characterize the density $1/\tilde{L}_1(\alpha)$ of the rare resonant local degrees of freedom in the system — see also the discussion at the end of Section 2.4.4.

At weaker disorder, $\alpha \in \{1, 1.2\}$ the naïve exponential fit would produce a very small $\tilde{L}_1(\alpha) < 2.5$. Moreover, the slope $d(\log N_E)/dL$ of the corresponding lines shows a clear increase as the system size grows. Accordingly, the fraction $N_E/D_{L,S,M}$ (bottom panel of Fig. 2.9) displays a tendency towards saturation, suggesting that, ultimately, the scaling of N_E in long systems becomes ergodic, $N_E \propto 2^L$.

In view of the results above, it may be tempting to conclude that the strongly disordered Heisenberg spin chains do indeed display a non-ergodic, non-MBL phase with unusual tree-like eigenstates that are only slightly dressed by perturbative corrections and occasional resonances (similar to how in the conventional MBL phase the eigenstates are perturbatively dressed product states). At weaker disorder, one would then expect a transition into an ergodic phase. However, the crucial question concerns the ultimate fate of the putative non-ergodic behavior in the thermodynamic limit. In particular, does the observed fractal scaling $N_E \propto 2^{L/\tilde{L}_1(\alpha)}$ persist, or does it eventually cross over to the ergodic scaling, as for the weakly disordered case? In order to answer these questions, we will now subject the hypothetical non-ergodic phase to several stringent tests.

Level statistics

Our main goal in this section is to further characterize the non-ergodic behavior found above and its dependence on the system size. We will employ the standard diagnostics of ergodicity and its breakdown: the level statistics in the center of the many-body band. We use a standard shift-and-invert technique in order to extract about 50 eigenstates per realization (fewer for $L \leq 12$ and $L = 26$) from the infinite-temperature region of the spectrum, namely, around the target energy

$$E_{\text{middle}} = \text{Tr} [\rho_{T=\infty} H] = \frac{\text{Tr} H}{D_{L,S,M}}, \quad (2.72)$$

which is easily computed for any given instance. At least 1 000 disorder realizations were sampled for most values of (L, α) .

From now on, unless otherwise specified, we are going to focus on the $S = M = 0$ sector of the Hilbert space and L even, with dimension $D_{L,0,0} = C_{L/2}$ (C_n being the n -th Catalan number). The data for $S \in \{1, 2\}$ did not show any qualitative differences.




Figure 2.10: Level statistics for the Heisenberg chain. Each curve in the figure was produced using at least 50 eigenstates from the middle of spectrum and 1000 disorder realizations. The dashed lines are the Wigner–Dyson (WD) and Poissonian distributions. (Left) For a fixed length $L = 22$ and varying α . (Right) For a fixed $\alpha = 1.3$ and varying length L . The tendency towards the WD statistics is evident both at growing L and α . However, for smaller values of α , $P(r)$ remains close to the Poisson one up to largest available system sizes.




Figure 2.11: Average r -parameter as a function of L for different values of the disorder α . From top to bottom (for any fixed L) $\alpha \in \{1.9, 1.6, 1.3, 1, 0.8, 0.6, 0.45, 0.3\}$. The dashed lines at $r \approx 0.53$ and $r \approx 0.39$ represent the WD and Poisson values, respectively. Error bars are within the symbol.

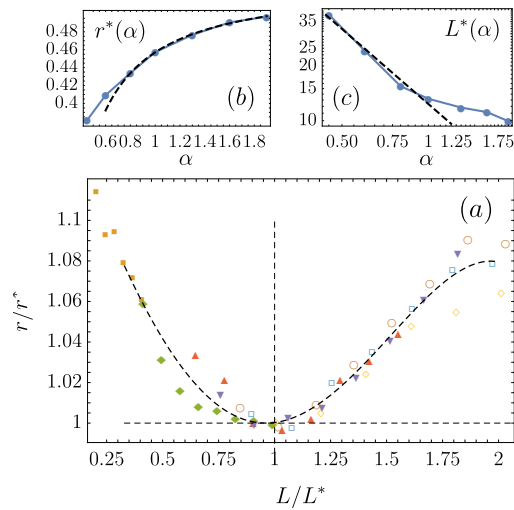


Figure 2.12: (a) The average r parameter as a function of L/L^* , divided by r^* , falls onto a universal curve in the vicinity of its minimum; (b) The value of the minimum $r^*(\alpha)$. The dashed line is a fit of the form $r^*(\alpha) = r_\infty + c_1/\alpha + c_2/\alpha^2$, which returns $r_\infty = 0.53(1)$ compatibly with the GOE value; (c) The position of the minimum $L^*(\alpha)$ for $\alpha \in \{1.9, 1.6, 1.3, 1, 0.8, 0.6, 0.45\}$. The dashed line is a fit of the form $L^*(\alpha) = c \alpha^{-\nu}$ with $\nu = 1.40(13)$ for the first 4 points.

We characterize the level statistics by the r -parameter, defined as in Eq. (2.7). Its distribution and dependence on the system size and disorder strength are shown in Figs. 2.10 and 2.11. The distribution of r_n changes qualitatively as α is decreased at a fixed L : for the largest $\alpha = 1.6$ (very weak disorder), it is described by the standard Wigner–Dyson distribution, while for small $\alpha = 0.3$ (strongest disorder considered) one observes the Poissonian distribution, with virtually no level repulsion. This supports the existence of a non-ergodic regime at accessible system sizes. For $\alpha \in [0.6, 1]$, the level statistics is intermediate between the Wigner–Dyson and Poissonian distributions. We also illustrate the dependence of the distribution $P(r)$ on the system size for weak disorder $\alpha = 1.3$. It is evident that the distribution flows towards Wigner–Dyson, albeit relatively slowly.

Further, we study the flow of the disorder- and eigenstate-averaged value, $\langle\langle r \rangle\rangle$,³⁴ with the system size, in an attempt to extract some relevant length scales. $\langle\langle r \rangle\rangle$ is illustrated in Fig. 2.11 as a function of L for different values of α . For weak disorder, $\alpha \geq 0.8$, the dependence of $\langle\langle r \rangle\rangle$ on L is not monotonic. Our data show a tendency towards the Poisson statistics for small system sizes, $L < L^*(\alpha)$, but for $L > L^*(\alpha)$ the value of $\langle\langle r \rangle\rangle$ starts growing, moving towards the Wigner–Dyson (WD) value. Upon decreasing α to the value of 0.8, the length scale $L^*(\alpha)$ increases, while its value $r^*(\alpha) = r[L^*(\alpha)]$ decreases. The ultimate flow of $\langle\langle r \rangle\rangle$ towards the WD value is consistent with the expectation that at weak disorder the system becomes ergodic for modest system sizes. One can estimate the scale where the system becomes ergodic, L_{erg} , by extrapolating the $\langle\langle r(L) \rangle\rangle$ dependence until the crossing with the WD line. The length scale extracted in this way is larger than the maximum system sizes accessible numerically for $\alpha < 1$. The extrapolation procedure suffers from a large uncertainty. Therefore, we choose instead to characterize the delocalization crossover by the length $L^*(\alpha)$, and we expect that $L_{\text{erg}}(\alpha) \propto L^*(\alpha)$.

The data at stronger disorder, $\alpha \in [0.3, 0.6]$, shows at least *prima facie* a qualitatively different behavior. For the strongest disorder, $\alpha = 0.3$, the parameter $\langle\langle r \rangle\rangle$ slowly increases for small L , in a stark contrast with the behavior found for $\alpha \geq 0.8$. Interestingly, at small L this parameter is below the Poisson value of $\langle\langle r \rangle\rangle_{\text{P}} \approx 0.39$. We attribute this to strong disorder leading to the appearance of very small couplings in a typical disorder realization (smaller than the level spacing at small L). The chain is then effectively broken into smaller, almost noninteracting, spin chains. This leads to level clustering and the r -parameter becomes sub-Poissonian. However, since the level spacing decreases exponentially with the system size, while the weakest coupling only decreases as a power law (see Eq. (2.65)), the level clustering is eventually washed out and for $L > 18$ the parameter $\langle\langle r \rangle\rangle$ rapidly approaches the standard Poisson value. For disorder strength $\alpha = 0.45, 0.6$, $\langle\langle r \rangle\rangle$ is initially slightly above the Poisson value, but it *decreases* as the system size is increased; no flow towards WD is seen. For the system sizes analyzed, it is evident that ergodicity has not developed and a single SDRG tree state provides a good approximation to the eigenstates,

³⁴With slight abuse of notation, we denote by $\langle \cdot \rangle$ the arithmetic mean over the roughly 50 eigenstates of a given disorder realization, which tends to the actual thermal average when $L \rightarrow \infty$ (notice, however, that r_n is not an observable).

as discussed previously.

The exact diagonalization results for strong disorder values, $\alpha \in [0.3, 0.6]$, may be consistent with two scenarios. One scenario is that (much like in the usual MBL) the system experiences a phase transition at some critical disorder strength. Another scenario is that, even at strong disorder, the system would eventually flow to ergodicity, similar to what we found for weaker disorder values. Assuming that this second scenario is realized, in large enough systems the curves for $\alpha \in \{0.45, 0.6\}$ would first develop a minimum and then flow to the WD value at yet larger system sizes. The corresponding scale L^* can be heuristically extracted by extrapolating³⁵ the ED data shown in Fig. 2.11. The dependence of the length L^* on disorder, as extracted by the analysis outlined above, is illustrated in Fig. 2.12(c). It is consistent with a power-law scaling, $L^*(\alpha) \propto \alpha^{-1.4}$. We note that the curves $\langle r(L) \rangle$ for $\alpha \geq 0.6$ (including extrapolated data for $\alpha \in \{0.45, 0.6\}$) can be roughly collapsed (in the vicinity of $L = L^*$) into a single one by simultaneously rescaling $r \rightarrow r/r^*$ and $L \rightarrow L/L^*$, as in Fig. 2.12(a).

To sum up, the length L^* beyond which the spectral parameter starts flowing towards the WD value (but the system of size L^* is still non-ergodic, because r^* is closer to the Poisson value), grows rapidly with the increase of disorder. Although the trend is clear, we are extrapolating significantly away from the accessible system sizes, $L \leq 26$. Thus, the law governing $L^*(\alpha)$ which we propose should be taken with a grain of salt. In the next part of our analysis, we proceed to test the eigenstate thermalization hypothesis.

Eigenstate thermalization hypothesis and its breakdown

We now characterize the eigenstates of random-bond Heisenberg chains by testing the the Eigenstate Thermalization Hypothesis (ETH) and its breakdown. The ETH was discussed in some detail in Section 1.1.2. For our current purposes, we focus on its formulation in terms of the expectation values of local observables \hat{O} , Eq. (1.18), which we repropose here for convenience (see the original discussion for an explanation of the various terms):

$$\langle a | \hat{O} | b \rangle = \bar{O}(E) \delta_{ab} + e^{-S(E)/2} f(E, \Delta E) R_{ab}. \quad (2.73)$$

We focus our attention on the following two local observables:

$$\hat{O}_{\max} = \mathbf{s}_{i^*} \cdot \mathbf{s}_{i^*+1}, \quad (2.74)$$

$$\hat{O}_{\text{rand}} = \mathbf{s}_{j^*} \cdot \mathbf{s}_{j^*+1}, \quad (2.75)$$

where $(i^*, i^* + 1)$ is the maximally-coupled spin pair,

$$|j^*| = \max_i |J_i| \quad (2.76)$$

³⁵The extrapolation is performed and the value of L^* is extracted by fitting the available data for $L \geq 14$ onto a quadratic function, $r(L) = r^* + a(L - L^*)^2$, with fitting parameters r^* , L^* and a .

and $(j^*, j^* + 1)$ its antipodal pair, $j^* = (i^* + L/2) \bmod L$. Since couplings are independent, the latter pair is coupled by an interaction J_{j^*} of typical (or “random”) strength, hence the name \hat{O}_{rand} .

Let us discuss our expectations for the averages of these operators over eigenstates, depending on whether SDRG is accurate. First, suppose that $|a\rangle$ is exactly an SDRG tree state. Then the spins $(i^*, i^* + 1)$ are going to be paired in either a $S = 0$ or a $S = 1$ state, and the value of $\langle a | \hat{O}_{\text{max}} | a \rangle$ is going to be either $-3/4$ or $1/4$, respectively. Even for $\langle a | \hat{O}_{\text{rand}} | a \rangle$, these two values are going to be likely, although in many cases the pair $(j^*, j^* + 1)$ will not be coupled directly by the SDRG procedure, but rather at a higher level, resulting in some intermediate value. However, in the ergodic regime — when SDRG breaks down — local thermalization implies that the local state of any pair of spins will be a uniform (at $T = \infty$) mixture of the four possible above-mentioned states, resulting in a thermal average of zero for both observables.

The distributions of the expectation values of $\hat{O}_{\text{max/rand}}$ over eigenstates at system size $L = 20$ are shown in Fig. 2.13. It is clear that the system is perfectly compliant with the ETH at sufficiently high values of α , whereas at smaller values of α the behavior is consistent with the eigenstates being close to tree states. This phenomenology, which we interpret as a finite-size crossover between ergodic and non-ergodic structure of the system’s eigenstates, is compatible with the behavior of the level statistics observed in the previous section.

In order to validate our interpretation, we characterize the finite-size flow to ergodicity by looking at the percentage of eigenstates whose corresponding values of $\hat{O}_{\text{max/rand}}$ falls within some fixed window centered at zero. Fig. 2.14 confirms that the “ergodic fraction” of infinite- T eigenstates is increasing with L for both \hat{O}_{max} and \hat{O}_{rand} , though much more slowly for strong disorder. In particular, at $\alpha = 0.6$ ETH is still strongly violated even at $L = 26$.

Entanglement entropy

Another witness of the non-ergodic behavior can be found in the scaling of the half-chain entanglement entropy with the system size, which is known to obey an area law for MBL systems, and a volume law for ergodic ones (cf. Section 2.2). More precisely, in a thermalizing system generic eigenstates are expected to be similar to random states; their entanglement entropy equals the thermodynamic one, yielding the expression valid for states in the middle of the band: $S_{\text{ent}}(L/2) = L/2 + o(L)$, when measured in bits [17, 197].

The numerical results are reported in Fig. 2.15. The median entanglement entropy of the infinite-temperature eigenstates exhibits linear scaling for all considered values of α , but the linear coefficient observed at $L \leq 20$ deviates substantially from the ergodic prediction at strong disorder (although significant curvature is present). This is once more consistent with the results of the foregoing sections.




Figure 2.13: Heat maps for the distributions of $\langle a | \hat{O}_{\max} | a \rangle$ (left) and $\langle a | \hat{O}_{\text{rand}} | a \rangle$ (right) over several ($\gtrsim 25\,000$) eigenstates. Here $S = 0$, $L = 22$ and $\alpha \in \{0.3, 0.6, 1.3\}$ (top to bottom). The concentration of the y -marginals around 0 denotes increasingly ergodic behavior (see the comments in the text).

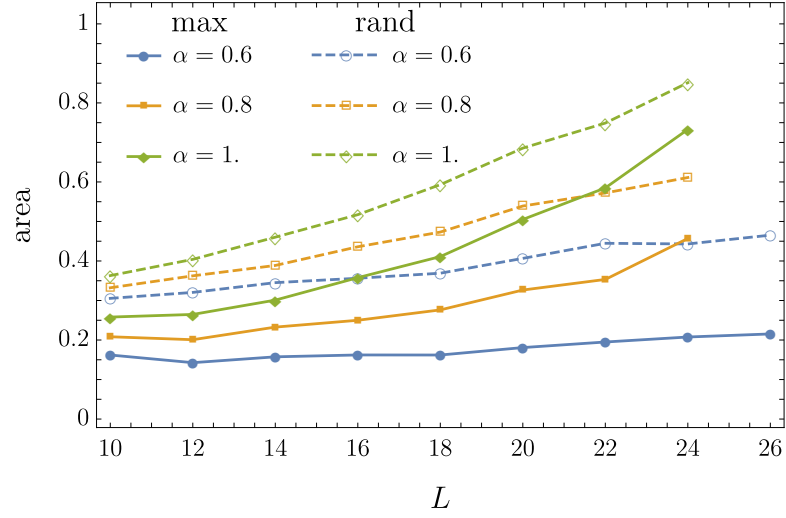


Figure 2.14: Fraction of eigenstates with a value of $\langle \hat{O}_{\max/\text{rand}} \rangle$ between $-1/8$ and $1/8$, for $\alpha \in \{0.6, 0.8, 1.0\}$ and $S = 0$. ETH predicts this fraction to become 1 in the infinite-size limit.




Figure 2.15: Median value of the half-chain entanglement entropy, $S_{\text{ent}}(L/2)$, for the $S = 0$ sector and $\alpha \in \{1.6, 1.3, 1, 0.8, 0.6, 0.45, 0.3\}$ (top to bottom). Linear fits, performed on the $L \geq 14$ points, are shown when their slope is close to the ergodic prediction. Entanglement entropy for the higher disorder values is strongly sub-thermal.

In summary, the ED data show a clear trend towards ETH for moderate-to-weak disorder (namely $\alpha \gtrsim 0.6$) while indicating a novel non-ergodic regime for the case of strong disorder. To determine the behavior of the system in the thermodynamic limit, which is totally inaccessible to ED, we have to resort to a completely different approach, as presented in the next section.

2.4.4 Resonance counting

As we showed in the previous section, at strong disorder finite-size random Heisenberg chains exhibit a non-ergodic regime in which their eigenstates are well approximated by tree states. Here, to determine the eventual fate of these systems in the thermodynamic limit $L \rightarrow \infty$, we develop an approach to analyze resonances between different tree states. We are able to capture long-range, multi-spin processes, which are beyond the scope of conventional SDRG. We obtain the asymptotic behavior of the resonance number, as well as their spatial structure. We will find that the resonance density grows for all studied disorder strengths, leading to an eventual delocalization at very large length scales, which we estimate. Beyond this length scale, the system presumably becomes ergodic.

Given a tree state generated by SDRG, we can construct a complete basis, Eq. (2.70), in the Hilbert space (with fixed total spin of the system) by allowing the values of the block spins identified by SDRG to take all possible values consistent with $\mathfrak{su}(2)$ fusions. The Hamiltonian (2.61), written in this basis, will then connect the initial SDRG state to a certain number of other tree states. We will consider the eigenvalue problem in this basis. Localization in this problem corresponds to true eigenstates being close to the tree states; in contrast, delocalization signals the breakdown of the SDRG approximation, suggesting ergodicity. The criteria for delocalization will be studied below.

Connectivity of the hopping problem

First, we investigate the connectivity of this eigenvalue problem. That is, we analyze how many matrix elements of the Hamiltonian between a given tree state and other ones are nonzero. The $SU(2)$ symmetry of the model imposes stringent constraints on the matrix elements of the Hamiltonian [212]. Specifically, let us consider one of the terms in the Hamiltonian, $J_i \mathbf{s}_i \cdot \mathbf{s}_{i+1}$. It can be shown that the action of such an operator on a tree state can only affect the block spins that lie on the fusion path connecting spins i and $i + 1$, see Fig. 2.6(c).

Moreover, each of the block spins on the path may only change by $\Delta S = 0$ or ± 1 . It then follows that the number of states connected to a given one by the operator $J_i \mathbf{s}_i \cdot \mathbf{s}_{i+1}$ is given by:

$$K_{i,i+1} \approx 3^{l_{i,i+1}} \quad (2.77)$$

where $l_{i,i+1}$ is the length of the fusion path of physical spins i and $i + 1$. The \approx sign is from the slight overcounting due to neglecting additional constraints (see Footnote 29 above in




Figure 2.16: Typical connectivity K of the Heisenberg Hamiltonian in the tree basis obtained from SDRG, as a function of system size, for different values of the disorder α (see legend). The dashed line represents the fit $K = (L/L_0)^\kappa$. The extracted value of $\kappa = 2.75$ is in good agreement with the result (2.81) for the random fusion ensemble, $\kappa = \log 3 / \log(3/2) \approx 2.71$. The horizontal dotted line shows the maximal sampling rate (2×10^7 matrix elements) used in the search of resonances (see Fig. 2.17). Above $L \approx 2 \times 10^3$, exhaustive search becomes too expensive and K_{res} is estimated via random search among connected matrix elements.




Figure 2.17: Average number of resonant neighbors for an SDRG tree state in short systems. The solid lines of different color correspond to different values of the disorder α . Dashed lines represent linear fits, $\langle K_{\text{res}} \rangle = aL - b$. For each value of α , the legend also indicates the scale $L_1(\alpha)$ defined by $\langle K_{\text{res}} \rangle = 1$ (see main text).

this chapter). Sufficiently far from the bottom of the tree, the typical values of the block spins are large and the extra constraints cannot alter the scaling in Eq. (2.77) by more than a prefactor. Taking into account that the Hamiltonian (2.61) is just a sum of local terms of the form discussed above, we conclude that the total connectivity in the Hilbert space induced by the Hamiltonian (2.61) is:

$$K \approx \sum_{i=1}^L 3^{l_{i,i+1}}. \quad (2.78)$$

We are now left with the task of computing the distribution $P(l)$ of these lengths $l_{i,i+1}$ for SDRG trees.

The SDRG fuses spins that are most strongly coupled. Neglecting the correlations between the (renormalized) couplings at any step of SDRG, as well as the dependence of those on the couplings at *earlier* stages of SDRG, we can assume that the pair of spins to be fused is just randomly chosen among all possibilities (with the only requirement that the fusing spins be nearest neighbors so that locality is respected). In such an ensemble of *random fusion trees* the distribution $P(l)$ can be computed analytically. As we show in Appendix C.1, it turns out that $P(l)$ decays exponentially with l , and in the limit $L \rightarrow \infty$ it becomes:

$$P(l) = \frac{3}{4} \left(\frac{2}{3} \right)^l \quad (2.79)$$

(the normalization is the correct one considering $l \geq 2$, so $\sum_{l \geq 2} P(l) = 1$). With this distribution $P(l)$, the sum (2.78) is dominated by the maximum l_M over the L terms. To leading order in L , the value of l_M is essentially deterministic and can be estimated from the condition $LP(l_M) \sim 1$.³⁶

This yields

$$l_M = \max_{i=1, \dots, L} l_{i,i+1} \sim \frac{\log L}{\log(3/2)}. \quad (2.80)$$

Plugging back into $K \sim 3^{l_M}$ we find

$$K \sim L^\kappa, \quad (2.81)$$

with $\kappa = \log 3 / \log(3/2) \approx 2.71$. The power-law scaling (2.81) and the value of the exponent κ are in a good agreement with the numerical simulations of the SDRG trees (with the full set of SDRG rules taken into account), see Fig. 2.16.

³⁶This follows from the distribution of the largest of L i.i.d. random variables with distribution Eq. (2.79), which is given by

$$P_M(x) = \frac{3L}{4} \left(\frac{2}{3} \right)^x \left(1 - \left(\frac{2}{3} \right)^{x-1} \right)^{L-1}.$$

Local resonances

Our next goal is to find resonances among the $K \propto L^k$ “hopping” processes generated by the Heisenberg Hamiltonian for a given tree state. We will first focus on investigating relatively small system sizes, $L \lesssim 30$, comparable to those accessible by exact diagonalization.

To study the number of resonances, we first use the SDRG procedure to generate a random tree state $|\Psi_{\text{RG}}^0\rangle$ and identify its resonant neighbors — that is, those states $|\Psi_{\text{RG}}^a\rangle$ for which the ratio

$$\left| \frac{\langle \Psi_{\text{RG}}^a | H | \Psi_{\text{RG}}^0 \rangle}{E_a - E_0} \right| \quad (2.82)$$

is larger than one.³⁷ These resonances invalidate the perturbative expansion around the infinite-disorder eigenstates (the SDRG states). Their proliferation signals the instability of tree states, strongly suggesting that ergodicity is restored. The SDRG is essentially a local optimization procedure that aims to construct basis states free of such resonances. Based on the results presented above, we expect that, at strong disorder and in relatively short systems, these resonances should be few in number, because SDRG is accurate.

The average number of resonant neighbors $\langle\langle K_{\text{res}} \rangle\rangle$ of an SDRG tree state is shown in Fig. 2.17. We observe that for relatively small systems this number scales linearly with the system size L . As expected, the slope of this linear growth becomes smaller for stronger disorder.

The condition $\langle\langle K_{\text{res}} \rangle\rangle = 1$ defines an important (disorder-dependent) length scale in the problem, $L_1(\alpha)$, at which resonances start appearing. Naïvely, this length scale plays the same role as the length scale L_W introduced in Section 2.4.1 to characterize the resonances in a random-field XXZ chain. We found that the scale $L_1(\alpha)$ grows at stronger disorder, crudely following a power-law dependence, $L_1(\alpha) \propto \alpha^{-0.4}$.

Fig. 2.17 shows that at relatively strong disorder values, $\alpha \leq 0.6$, the average number of resonant neighbors for an SDRG tree state is 1 or less for all system sizes available in ED. This agrees with the observation that such chains display a non-ergodic behavior in all of the ED studies of the previous sections, with eigenstates being well approximated by tree states.

In particular, the low number of resonances is in agreement with the slow growth of N_E (the participation ratio of eigenstates in the tree basis) discussed below Eq. (2.71). Drawing parallels to the conventional MBL systems, it is tempting to identify the length scale $\tilde{L}_1(\alpha)$ that controls the exponential growth of N_E with the system size with $L_1(\alpha)$. However, the comparison of the values of $L_1(\alpha)$ and $\tilde{L}_1(\alpha)$ reveals that the latter is several times shorter. We attribute this difference to the effect of the second-order perturbative corrections that contribute to the spreading of the exact eigenstate $|E\rangle$ over SDRG tree state. Such higher-order contributions lie beyond the first-order resonance counting that underlies the scale

³⁷This widespread localization criterion is on the same line as the one proposed by Edwards and Thouless for the Anderson model [88].

$L_1(\alpha)$. The corrections are expected to be more significant at weak disorder; in accordance with this intuition, we found a more significant difference between $L_1(\alpha)$ and $\tilde{L}_1(\alpha)$ for such disorder strengths.

Longer systems and proliferation of resonances

Does the linear scaling of $\langle\langle K_{\text{res}} \rangle\rangle$ with the system size discussed in the previous section persist in the thermodynamic limit? Such behavior would bear a close resemblance to the strongly disordered XXZ model. It would imply that the resonant neighbors can be attributed to the existence of local subsystems with resonating levels which, if sufficiently separated in space, would remain isolated and not cross-talk (in the sense that no significant entanglement would exist in the eigenstates between such “local” resonances). If true, this would be a strong argument in favor of the SDRG tree states surviving in an infinitely long system, up to dressing corrections from local and isolated resonances. We now perform a detailed resonance analysis in large systems, up to $L \approx 2 \times 10^3$, finding that Heisenberg chains actually behave qualitatively differently compared to plain-vanilla MBL systems: the number of resonances grows *faster* than linearly with the system size.

The probability for an SDRG tree state to have no resonant neighbors vanishes in sufficiently long systems (see top panel in Fig. 2.18). Past that length scale, a typical tree state has a large number of resonances attached to it. The bottom panel in Fig 2.18 shows (in log–log scale) the dependence of the typical number of resonant neighbors (in the sense of $\exp(\langle\log K_{\text{res}}\rangle)$) for an SDRG tree state. The dashed lines represent power-law fits of the form

$$\exp(\langle\log K_{\text{res}}\rangle) \propto \left(\frac{L}{L_1(\alpha)} \right)^{1+\mu}, \quad (2.83)$$

with the scale $L_1(\alpha)$ determined from the short-scale behavior of K_{res} , see the previous section. The anomalous exponent μ is approximately disorder-independent, $\mu \approx 0.37$. The details of the numerical procedure employed to find and characterize resonances are given in Appendix C of Ref. [213].

The power-law scaling of the number of resonant neighbors, $K_{\text{res}} \propto L^{1+\mu}$, implies that, in stark contrast to the random-field XXZ model, the *density* of the resonating degrees of freedom grows with the system size. Accordingly, at least some of the resonant transitions must originate not from the rearrangement of a few local spins, but rather involve a growing number of spins. To support this conclusion, we analyze the structure of typical resonances. The left panel of Fig. 2.19 shows the average number of block spins, N_{bs} , that are changed in the course of a resonant transition. We observe that N_{bs} grows (albeit rather slowly) with the length of the chain. The decrease of N_{bs} with increasing disorder can be understood as follows: at weak disorder, the possibility to change a block spin in a resonant manner is often accompanied by an “instability” (with respect to resonances) of the block spins higher up in the hierarchy. Therefore, increasingly nonlocal resonant neighbors appear and contribute to the increase of the average N_{bs} . On the other hand, at

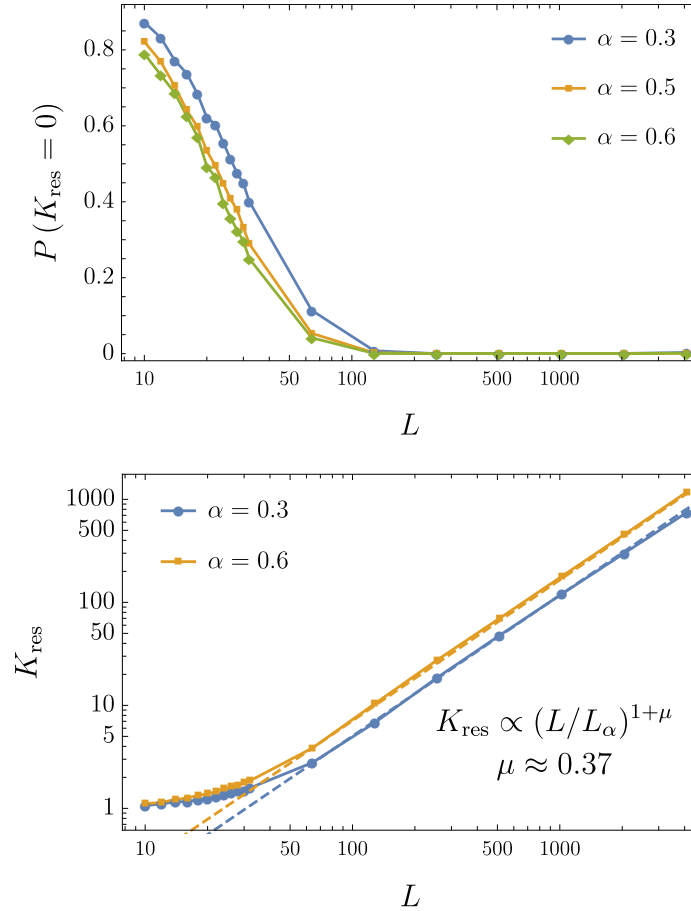


Figure 2.18: (Top) The probability $P(K_{\text{res}} = 0)$ for an SDRG tree states to have no resonant neighbors, as a function of the system size, and for different values of disorder α (see legend). Note the logarithmic scale along the horizontal axis. While short systems are essentially free of resonances even for a relatively weak disorder, the probability $P(K_{\text{res}} = 0)$ becomes vanishingly small in long systems. (Bottom) Typical number of resonant neighbors for an SDRG tree state in long systems, $L \leq 2^{12}$. The solid lines correspond to the numerical results for $\alpha = 0.3, 0.6$, while the dashed lines are power-law fits, $K_{\text{res}} = [L/L_1(\alpha)]^{1+\mu}$. The scale $L_1(\alpha)$ where the average number of resonant neighbors for an SDRG tree state equals 1 was defined previously in the main text.

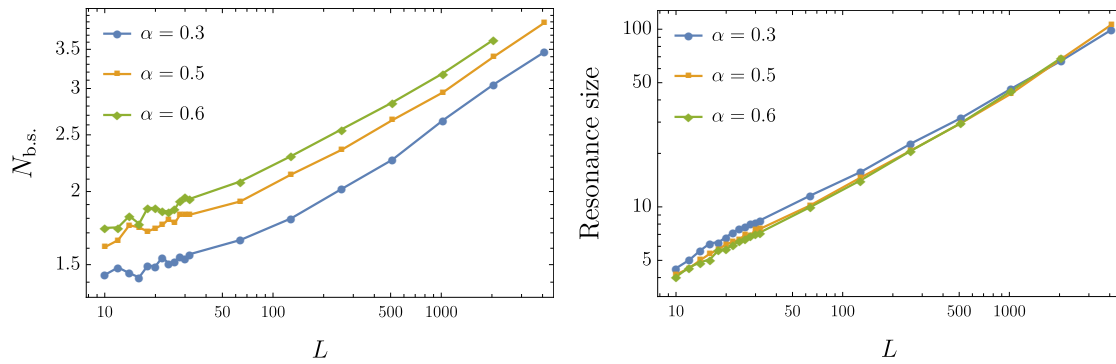


Figure 2.19: (Left) Average number of the block spins changing their value in a single resonant transition. (Right) Typical number of adjacent physical spins involved in a resonance (see main text).




Figure 2.20: Characteristic energy scale for typical resonances, as a function of system size, shown for different disorder strengths. See main text for the definition of V_{typ} .

strong disorder an “instability” of a single block spin is more likely to remain “localized” and not to “propagate” upwards in the tree, resulting in a smaller N_{bs} .

The size of a typical resonance in *real space* also grows as the system size is increased (see the right panel of Fig. 2.19). A physical spin is affected by a resonant transition if one or more spins on a fusion path rooted at that physical spin changes its state. The physical size of a resonance is then defined as the total number of leaves within the causal cone below the highest changed spin. A more detailed explanation is provided in Appendix C.2, as well as a proof that the resonance size is $O(\log L)$ in the random fusion ensemble.

We observe that at moderate system sizes the typical spatial size of a resonance at strong disorder exceeds that at weak disorder. This is in accord with our intuition: at strong disorder a large number of spins need to rearrange collectively in order for a transition to be resonant. In terms of the SDRG, this means that many SDRG steps can be performed before the resonances start to play any role. On the other hand, in sufficiently long systems we see the opposite tendency: weakly disordered chains typically exhibit resonances of larger size. This is the manifestation of the propagation of an “instability” of block spins upwards in the tree, as discussed above.

The growing length scale characterizing the resonances is accompanied by a corresponding (decreasing) energy scale, namely the typical matrix element V_{typ} for a resonant transition. Its system size dependence is shown in Fig. 2.20. In the following and final part of our investigation, we are going to use V_{typ} to estimate the energy scale associated with the crossover to ergodicity.

Breakdown of SDRG and delocalization

The results presented in the previous section (most importantly, the power-law growth of the resonance density) strongly suggest that even in the strongly disordered chains with $\alpha \leq 0.6$, where the ED studies of Section 2.4.3 reveal little (if any) signs of ergodicity, the resonant transitions missed by SDRG eventually proliferate. In this section, we estimate the corresponding thermalization scale $L_{\text{erg}}(\alpha)$.

Given an SDRG tree state and a set of resonant transitions associated with it, one can identify a set of block spins that can be changed via at least one resonant process. We refer to those block spins as resonant, or unstable ones. For a chain of L spins there are $2L - 1$ nodes in the SDRG tree (L physical spins plus $L - 1$ block spins). At each stage of the SDRG procedure some number L_{RG} of the block spins play the role of effective “physical spins” of the system. For example, in the initial state of the SDRG we have $L_{\text{RG}} = L$ and the SDRG spins are just the physical ones. The final stage corresponds to $L_{\text{RG}} = 1$, with the top node of the SDRG tree being the only remaining spin. The ratio L/L_{RG} is nothing but the average size of the spin clusters in the system.

At any given moment during the SDRG, only the unstable block spins that are among the L_{RG} spins currently comprising the system are relevant for potential delocalization. The others either have not formed yet, or have already been decimated; they are not

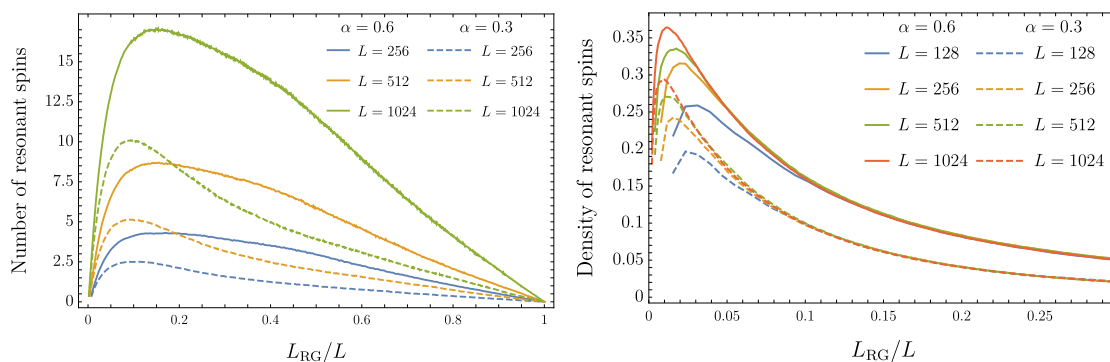


Figure 2.21: Evolution of the number (left) and density (right) of the resonant blocks spins in the course of SDRG. The horizontal axes shows the ratio of the current system size to the initial length of the system (so the RG runs “leftwards”).

expected to contribute directly to the physics at the current energy scale. It is thus natural to ask how the number and density of the resonant block spins evolve over the course of the renormalization process.

This is illustrated in Fig. 2.21, that shows the dependence of the number (left panel) and the density (right panel) of unstable block spins for two different values of α and several values of the physical chain length L . These quantities are plotted as a function of the running RG length L_{RG} , normalized by L . For not too small L_{RG}/L , we observe that for a fixed disorder strength the density of resonant spins exhibits a universal (L -independent) behavior, $\rho_{\text{res}}(L_{RG}/L, L) = \rho_{\text{res}}(L_{RG}/L)$. The density $\rho_{\text{res}}(L_{RG}/L)$ is higher at weaker disorder and grows throughout the SDRG. In contrast, at small L_{RG}/L (corresponding to the final stages of SDRG in a finite chain) a rather pronounced dependence on L is observed.³⁸

Next, let us denote by $\rho_{\text{max}} = \rho_{\text{max}}(\alpha, L)$ the maximum density of unstable spins developed during SDRG process. Small ρ_{max} means that ρ_{res} remains small at all steps of SDRG. We then expect the resonances to be of little importance for our system. On the contrary, $\rho_{\text{max}} \sim 1$ indicates that at some stage the SDRG inevitably runs into a state where almost all spins participate in resonances. Then, the basic assumptions of SDRG are violated and we expect it to break down — this means that block spins are no longer well-defined, and start resonating. Presumably, this signals the onset of ergodicity.

It is natural to assume that there exists a critical value $\rho_{\text{max},c} < 1$ at which the SDRG falls apart, *i.e.* where the system crosses over from non-ergodic to ergodic behavior. We

³⁸Notice that this size-dependent drop in the resonance density should be expected, as the decimation must eventually overtake the resonance proliferation in a finite chain. In the thermodynamic limit, one can expect (and it is indeed suggested by the universal behavior of the curves in Fig. 2.21) that the density of resonances will continue to grow for arbitrarily small values of L_{RG}/L .

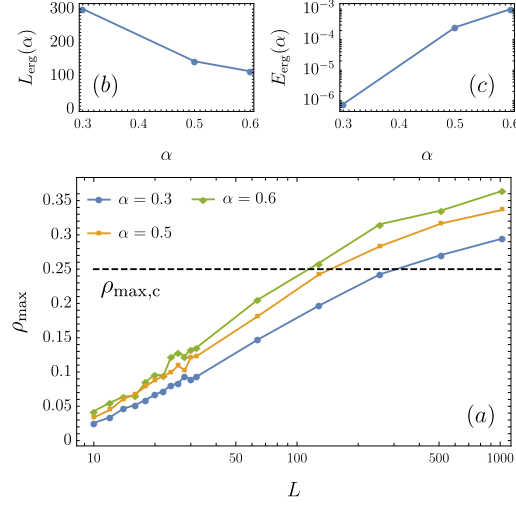


Figure 2.22: (a) Maximum density ρ_{\max} of resonant spins developed during the SDRG (cf. Fig. 2.21) as a function of system size. The condition $\rho_{\max}(\alpha, L) = \rho_{\max,c}$ defines the ergodization length scale $L_{\text{erg}}(\alpha)$ and the corresponding energy scale $E_{\text{erg}}(\alpha) = V_{\text{typ}} [L_{\text{erg}}(\alpha)]$. (b)–(c) Estimates for the ergodization length and energy scales obtained by fixing the critical density at $\rho_{\max,c} = 0.25$.

can then identify the length of the system at which $\rho_{\max} = \rho_{\max,c}$ with the ergodicity scale:

$$\rho_{\max} [\alpha, L_{\text{erg}}(\alpha)] = \rho_{\max,c}. \quad (2.84)$$

The scale $L_{\text{erg}}(\alpha)$, along with the corresponding typical matrix element for resonant transitions V_{typ} , gives an estimate for the ergodicity time and energy scales:

$$\tau_{\text{erg}} = E_{\text{erg}}^{-1}, \quad E_{\text{erg}} = V_{\text{typ}}(L_{\text{erg}}). \quad (2.85)$$

Figure 2.22 shows the dependence of $\rho_{\max}(L)$ for different disorder strengths. Estimating $L_{\text{erg}}(\alpha)$ requires fixing the critical density $\rho_{\max,c}$. While we have no general theory for $\rho_{\max,c}$, we observe (see Fig. 2.22) that a value $\rho_{\max,c} \in [0.2, 0.25]$ (similar to the critical density of resonances in the random XXZ model) results in an estimate $50 \lesssim L_{\text{erg}}(\alpha = 0.6) \lesssim 100$ that is essentially consistent with the intuition developed in the ED studies of Section 2.4.3, with

$$L_{\text{erg}}(\alpha = 0.6) \approx 2L^*(\alpha = 0.6) \approx 50.$$

Thus, for the purpose of an estimate, we choose $\rho_{\max,c} = 0.25$. In any event, we find it unlikely that a different choice of $\rho_{\max,c}$ would alter the asymptotic scaling of the estimated L_{erg} , provided $0 < \rho_{\max,c} < 1$. The resulting values for the length and energy scales at which thermalization starts to occur are shown in Fig. 2.22.

It is evident that at strong disorder, $\alpha = 0.3$, resonances start to proliferate only at very large length scales $L_{\text{erg}} \approx 300$ and ridiculously long time scales, far beyond the limitations of present-day synthetic platforms, where ergodicity and its breakdown are actively investigated (see Ref. [2] for a review). Thus, in experiments, strongly disordered, SU(2)-symmetric systems are expected to display the novel non-ergodic regime described above.

Systems of size $L \gg L_{\text{erg}}$ will be slowly thermalizing, and presumably exhibit slow diffusive transport at low frequencies. An interesting open question concerns the eventual fate of the integrals of motion obtained in the first steps of the SDRG (when the typical cluster size is much smaller than L_{erg}). Such nearly-conserved operators arise due to strongly coupled clusters of spins, and therefore destroying them would typically involve a relaxation process with a large characteristic energy ΔE . In very large systems, slow thermalizing processes will eventually destroy the conservation of these operators. However, since thermalization processes typically occur on a much smaller energy scale, $E_{\text{erg}} \ll \Delta E$, we expect that the decay time of such operators will be parametrically large in $\Delta E/E_{\text{erg}}$. An instructive example is that of a narrow-bandwidth thermal bath with energy scale E_0 ; there, the relaxation of excitations with energy $\omega \gg E_0$ is exponentially slow in ω/E_0 [3]. We expect that the integrals of motion obtained within SDRG before its breakdown will be similarly long-lived. Thus, our proposed picture is that the dynamical properties of strongly disordered Heisenberg chains are going to be captured by SDRG at frequencies $\omega \gtrsim E_{\text{erg}}$ (in particular, they would have nontrivial noise properties, described in Ref. [8]). At lower frequencies, $\omega \lesssim E_{\text{erg}}$, the system will undergo a crossover to regular diffusive behavior.

2.4.5 Conclusions

To sum up, the goal of this Section 2.4 was to investigate the effects of continuous non-Abelian symmetries on dynamical properties of disordered systems. We have considered a concrete example of disordered Heisenberg spin chains, with eigenstates characterized by an SU(2) symmetry. To describe the properties of this model, we combined state-of-the-art exact diagonalization studies with a new approach that allows us to include long-range resonances into the strong-disorder renormalization group.

We have found that in a broad range of disorder strengths and system sizes, Heisenberg chains exhibit a new kind of non-ergodic behavior. In this regime, the highly excited eigenstates have a scaling of entanglement entropy that is intermediate between the area law characteristic of MBL states and the volume law found in thermalizing systems. This behavior stems from the tree tensor network structure of the eigenstates obtained within SDRG. Simultaneously, in this regime the system exhibits a different kind of integrability, with integrals of motions having a varying degree of locality: some of them act on a small number of neighboring spins, while others act on larger and larger spin clusters.

Furthermore, we found that for weak disorder the behavior crosses over from non-

ergodic to ergodic as the system size is increased. For stronger disorder, all system sizes accessible numerically exhibited non-ergodic behavior. To address the eventual fate of the non-ergodic phase in this case, we have extended the SDRG approach, characterizing resonances that endanger the stability of tree states. Our results strongly suggest eventual delocalization and ergodicity, albeit at very large system sizes; delocalization occurs via unconventional, multi-spin processes, which is yet another unique feature of disordered systems with non-Abelian symmetries. A possible line of future research is to describe the transition between non-ergodic and ergodic regimes as a function of system size. A promising starting point seems to be to formulate an effective model in terms of resonant degrees of freedom, with parameters extracted using the methods described above.

Another interesting direction is to better understand dynamical signatures of the new non-ergodic regime uncovered here. One natural experiment would be to probe the dynamics of the most local integrals of motion (*e.g.* the total spin of a pair of strongly coupled physical spins), and to observe that, for system sizes $L < L_{\text{erg}}$, it is conserved to a good precision and for arbitrarily large times. Another interesting open question concerns spin transport in N -species disordered Hubbard models [210, 224]. In case of a $SU(N = 2)$ flavor symmetry, our work suggests that a sufficiently large system should show thermalizing behavior. Further work is required to establish the details of the dynamics (*e.g.* diffusion vs. subdiffusion).

More broadly, this work sets the stage for future discovery of new non-ergodic regimes and true dynamical phases that survive in thermodynamic limit. The approach introduced here can be naturally extended to other symmetry groups, for example $SU(N)$ spins for any $N > 2$. Even more generally, it would be interesting to investigate the stability of other tree tensor network structures with intermediate entanglement scaling, as possible good approximation of eigenstates in physical systems.

Chapter 3

Simulating Quantum Tunneling

If we cannot pass over the mountain, let's go under it.

The Lord of the Rings: The Fellowship of the Ring
Screenplay by P. Jackson. 2001

3.1 Quantum annealing and energy barriers

In Section 1.3.2, we described the Quantum Annealing (QA) algorithm in some detail. In particular, it was pointed out that one of the main reasons why QA may be expected to outperform classical algorithms on certain kinds of problems is its ability to make use of tunneling, a quantum effect with no analogue in classical physics. In turn, this suggests that the classes of problems where QA should be expected to shine are precisely those characterized by energy landscapes featuring an abundance of low-lying states separated by large barriers hindering classical local search. It was also noted that spin glasses represent the perfect candidate as they present these very characteristics.

While the general idea seems attractive enough, it is not so straightforward to actually demonstrate or even convincingly test a model for quantum speedup. This is a general truism in classical-vs-quantum comparisons, as they either rely on exact diagonalization, which is by nature very limited in the maximum system size one can probe, or they have to concede some ground in accepting poorly-controlled approximation schemes or rather noisy data. The latter case is the more interesting for our current discussion, whereas the former was encountered in Chapter 2 and will appear again in Chapter 4.

In particular, we are going to devote this chapter to the effort of using Quantum Monte Carlo (QMC) methods — as defined in the following sections — to benchmark the QA algorithm. In order to set up a fair comparison, it is fitting to only consider *general purpose*

classical algorithm, *i.e.* ones that are not specifically tailored for the problem at hand, just like vanilla QA is general-purpose and makes no explicit use of properties of the problem Hamiltonian. For this reason, we make no claims of optimality when discussing the application of these algorithms to the specific problems we are going to analyze, and one should in fact expect custom algorithms to exist that outperform both QA and QMC in finding (or approximating) the ground-state energy of the models under examination. That being said, if such amazing algorithms were so easy to devise there would be no need for this discussion!

3.2 Quantum Monte Carlo methods

The term “Monte Carlo” encompasses a number of different classical methods whose common ingredient is stochasticity.¹ They typically operate by setting up a temporal chain of random variables with update rules which are probabilistic in nature, and so chosen as to guarantee that the distribution of the random variables converge to a desirable ensemble (*e.g.* the thermal distribution for a given inverse temperature β). In this way, the value of observables and other physical quantities of interest may be estimated by repeatedly sampling the system after this convergence has occurred. It is worth pointing out that at variance with molecular dynamics, where one directly integrates the differential equations describing the system’s evolution, in the Monte Carlo realm there is no obvious correspondence between the number of simulation steps and the physical time coordinate.

The main advantage of Monte Carlo methods is that they possess a favorable scaling of complexity with respect to the system’s dimension: roughly speaking, as long as efficient update rules can be found the precision of the algorithm will always scale like $O(1/\sqrt{N_s})$, with N_s the number of (sufficiently uncorrelated and equilibrated) samples, regardless of the dimensionality d , whereas molecular dynamics will require $\Theta(1/\varepsilon^d)$ grid points in order to achieve a fixed precision ε .

“Quantum” Monte Carlo (QMC) is simply the application of Monte Carlo methods to quantum systems, which can be accomplished in a variety of ways. We describe two families of QMC methods which are used ubiquitously in computational physics: the Path-integral Monte Carlo (PIMC) and the Diffusion Monte Carlo (DMC) algorithms [248].

3.2.1 Path-integral Monte Carlo

The PIMC algorithm is based on a standard quantum-to-classical mapping which turns the system’s Hamiltonian into a classical one, amenable to the standard Monte Carlo toolkit. This is accomplished via a Suzuki–Trotter expansion, which effectively trades the “quantumness” of the problem in exchange for an extra effective dimension. More

¹The name was coined by Metropolis in reference to the Monte Carlo casino in Monaco, another well-known den of stochasticity. See Ref. [180] for a historical overview.

precisely, let the system be described by a quantum Hamiltonian of the form $H = H_{\text{cl}} + H_{\text{qu}}$, with the classical part diagonal on some computational basis,

$$\langle \tilde{z} | H_{\text{cl}} | z \rangle = E(z) \delta_{\tilde{z}, z} \quad (3.1)$$

and H_{qu} an off-diagonal term.

The goal of PIMC is to sample the equilibrium properties of the system by estimating the partition function

$$Z(\beta) = \sum_z \langle z | e^{-\beta H} | z \rangle \quad (3.2)$$

or some similar object involving the computational-basis propagator.

By using the Lie–Suzuki–Trotter formula for two noncommuting operators A and B ,²

$$e^{A+B} = \left(e^{A/K} e^{B/K} \right)^K + O\left(\frac{1}{K}\right), \quad (3.3)$$

and then interfixing $K - 1$ resolutions of the identity,

$$\begin{aligned} \langle z_K | U_{K-1} \cdots U_1 U_0 | z_0 \rangle \\ = \langle z_K | U_{K-1} \left(\sum_{z_{K-1}} |z_{K-1}\rangle \langle z_{K-1}| \right) \cdots U_1 \left(\sum_{z_1} |z_1\rangle \langle z_1| \right) U_0 | z_0 \rangle, \end{aligned} \quad (3.4)$$

the imaginary-time propagator can be expanded as follows:

$$\begin{aligned} \langle \tilde{z} | e^{-\beta H} | z \rangle &= \sum_z \prod_{k=0}^{K-1} \langle z_{k+1} | e^{-\beta H/K} | z_k \rangle \\ &\approx \sum_z e^{-\beta \sum_{k=0}^{K-1} \frac{E(z_k)}{K}} \prod_{k=0}^{K-1} \langle z_{k+1} | e^{-\beta H_{\text{qu}}/K} | z_k \rangle \\ &= \sum_z e^{-\beta H_{\text{eff}}(z)}, \end{aligned} \quad (3.5)$$

where $z_0 = z, z_K = \tilde{z}, \mathbf{z} = (z_1, \dots, z_{K-1})$ and we assumed that it is possible to cast the infinitesimal off-diagonal propagator (*transfer matrix*) in the form

$$\langle z_{k+1} | e^{-\beta H_{\text{qu}}/K} | z_k \rangle = e^{-\beta T(z_k, z_{k+1})/K} \quad (3.6)$$

²More precisely, for any pair of operators (A, B) and any $K \in \mathbb{N}$ we have

$$\left\| e^{A+B} - \left(e^{A/K} e^{B/K} \right)^K \right\| \leq \frac{\|[A, B]\|}{2K} e^{\|A\| + \|B\|},$$

where $\|\cdot\|$ is any operator norm [244].

so that it may be recombined with the classical part to define the effective (classical) Hamiltonian

$$H_{\text{eff}}(\mathbf{z}) = \frac{1}{K} \sum_{k=0}^{K-1} \left[E(z_k) + T(z_k, z_{k+1}) \right]. \quad (3.7)$$

The index k labels imaginary time slices which are coupled together by the transverse term T .

Observation. Notice that the procedure outlined above is the same one typically used to derive the path-integral formulation of quantum mechanics. In that case, one usually has in mind a system of n particles moving in the d -dimension Euclidean space \mathbb{E}_d and subject to some (e.g. 2-body) potential V . In other words, one takes

$$H_{\text{cl}} = \sum_{i < i'} V(x_i, x_{i+1}), \quad H_{\text{qu}} = \sum_{i=1}^n \frac{p_i^2}{2m_i}$$

(in first quantization formalism, with x_i and p_i the canonical position and momentum operators for the i -th particle). After Trotterization, one is able to write the effective Hamiltonian as

$$H_{\text{eff}} = \frac{1}{K} \sum_{k=0}^{K-1} \left[\sum_{i=1}^n \left(\frac{x_{i,k} - x_{i,k+1}}{\beta/K} \right)^2 + \sum_{i < i'} V(x_{i,k}, x_{i',k}) \right] \quad (3.8)$$

(where $x_{i,k}$ is the position of the i -th particle in the k -th time slice) and perform the $K \rightarrow \infty$ limit to obtain the standard path-integral representation of the imaginary-time propagator:

$$\langle \vec{x}' | e^{-\beta H} | \vec{x} \rangle = \int \mathcal{D}\vec{\xi}(\tau) e^{-\mathcal{S}[\vec{\xi}(\tau)]}, \quad (3.9)$$

where $\vec{x} = (x_1, \dots, x_n)$ and $\vec{x}' = (x'_1, \dots, x'_n)$ describe the initial and final states in terms of well-defined particle positions (here $x_1, \text{etc.}$ are numbers, not operators), the integral is performed over the space of paths in configuration space $\vec{\xi} : [0, \beta] \rightarrow (\mathbb{E}_d)^n$ with $\vec{\xi}(0) = \vec{x}$, $\vec{\xi}(\beta) = \vec{x}'$ and

$$\mathcal{S}[\vec{\xi}(\tau)] = \int_0^\beta d\tau \left[\sum_{i=1}^n \frac{1}{2m_i} \left(\frac{d\xi_i}{d\tau} \right)^2 + \sum_{i < i'} V(\xi_i(\tau), \xi_{i'}(\tau)) \right] \quad (3.10)$$

is the classical action of the n -particle system. This explains the terminology “path-integral” Monte Carlo.

Coming back to our PIMC algorithm, we can see that the original problem has been reduced to an equivalent classical problem up to a “Trotter error” $O(\beta^2/K)$. Notice, in particular, that the number of time slices K needed to achieve a certain precision depends on the value of β , or from a more pictorial point of view, the inverse temperature determines the “length” one has to ascribe to the imaginary time dimension for a fixed Trotter precision. For this reason, the PIMC algorithm described above only works at finite temperature, as

for $T = 0$ there is no good way of discretizing the infinite extra dimension. That being said, ground-state properties can still be studied in gapped systems by adjusting the temperature to be appreciably smaller than the gap, in which case the contribution from the excited states will be negligible.

The partition function Eq. (3.2) is obtained from Eq. (3.5) simply by setting $z_0 = z_K$ and summing over it, which amounts to switching from open to periodic boundary conditions for H_{eff} . The summands $e^{-\beta H_{\text{eff}}(\mathbf{z})}$, which are positive, are then interpreted as (non-normalized) Gibbs weights for the paths $\{\mathbf{z}\}$.

Take notice of the crucial assumption contained in Eq. (3.6): as this condition needs to hold for any (z_k, z_{k+1}) , and in particular for $z_k \neq z_{k+1}$, and as we must assume $\beta/K \ll 1$ in order to have a small Trotter error, we can expand the left-hand side and see that the T matrix is only well-defined if

$$\langle \tilde{z} | H_{\text{qu}} | z \rangle \leq 0 \quad \text{for all } \tilde{z} \neq z. \quad (3.11)$$

If *any* off-diagonal matrix element of the quantum term is positive, we cannot proceed with the Suzuki–Trotter expansion: we have a *sign problem*. In some cases, we may be able to find a clever basis change so that condition (3.11) becomes fulfilled for the new basis. However, this is not possible in general. Moreover, even when such a basis change is available, it is only fruitful if the new basis is simple enough to be used as a new computational basis (e.g. it should be a product basis), otherwise the task of computing the weights $e^{-\beta H_{\text{eff}}(\mathbf{z})}$ might incur a prohibitive slowdown.

In the event that Eq. (3.11) holds for some known, simple basis $\{|z\rangle\}$, Hamiltonian $H = H_{\text{cl}} + H_{\text{qu}}$ is said to be *stoquastic* [56]. In this case, it is computationally easy to evaluate the Gibbs weights (which are guaranteed to be nonnegative) and hence implement the usual Metropolis algorithm [181] to reach the equilibrium distribution.

An important property of stoquastic Hamiltonians is that their ground state wave function can always be taken to be real and nonnegative [56, 57], as follows from the Perron–Frobenius theorem [40]. One is naturally lead to wonder whether stoquastic ground states, which avoid the sign problem as well as other interference effects, may always be efficiently approximated by classical algorithms. Although this seems to be true in several cases, to the point that it can well-nigh be granted the status of “folk theorem” among the community, it is known that it does *not* in fact hold in full generality. For one, Hastings showed how PIMC with local updates can fail to efficiently simulate systems living on topologically nontrivial spaces, leading to exponential slowdowns even in the presence of polynomial gaps [119]. This proves that QA is strictly superior to PIMC at least on this class of problems. As we will show in Section 3.3, topological obstructions hampering PIMC may be eluded by the diffusion Monte Carlo algorithm presented below. Nonetheless, examples of cost Hamiltonians which are solvable in polynomial time by QA but take exponential time with DMC were also explicitly constructed [58, 135], showing that not even DMC is capable of fully simulating adiabatic stoquastic dynamics. We will

further discuss the value of classically simulating quantum annealing in the introduction of Section 3.3, as well as in the general conclusions to this chapter.

3.2.2 Diffusion Monte Carlo

The DMC algorithm belongs to the family of so-called *projective* (PQMC) methods, *i.e.* it aims at estimating the ground-state wave function of a system by evolving a generic quantum state $|\psi\rangle$ in imaginary time and thereby projecting out its ground component:

$$\begin{aligned} e^{-\tau H} |\psi\rangle &= \sum_{a=0}^{D-1} e^{-\tau E_a} c_a |E_a\rangle = e^{-\tau E_0} c_0 |E_0\rangle \left(1 + \frac{c_1}{c_0} e^{-\tau(E_1-E_0)} |E_1\rangle + \dots \right) \\ &\sim e^{-\tau E_0} c_0 |E_0\rangle \quad \text{for } \tau \rightarrow \infty, \end{aligned} \quad (3.12)$$

where we assumed that the state $|\psi\rangle = \sum_a c_a |E_a\rangle$ has a nonzero projection $c_0 = \langle E_0 | \psi \rangle$ on the ground state $|E_0\rangle$, and there is no ground state degeneracy, $E_1 > E_0$.

Notice that we must take some care with this procedure in order to obtain meaningful results. Indeed, the imaginary time evolution operator is not unitary,

$$(e^{-\tau H})^\dagger \neq (e^{-\tau H})^{-1}, \quad (3.13)$$

and therefore it is not bound to preserve norms. A naïve application of Eq. (3.12) would generally lead to an exponential decay (or increment) in the norm of $|\psi\rangle$, which would render it useless for computation. For this reason, a renormalization procedure will have to be included in the evolution process.

Other than this detail, we can use the same idea as PIMC in order to approximate the evolution operator on a classical machine: discretize the (imaginary) time interval in K steps, use the Trotter formula to factorize the small-time evolution

$$e^{-\tau H/K} \approx e^{-\tau H_{\text{cl}}/K} e^{-\tau H_{\text{qu}}/K} \quad (3.14)$$

and inject resolutions of the identity between every step. This will again lead to a twofold evolution step, where the quantum part provides a kinetic term and the classical part a potential. We first describe the simplest scenario — a one-particle system in a fixed external potential — before switching to the many-body case.

A single particle can be described by its amplitude distribution in real space, namely its wave function $\psi(x) = \langle x | \psi \rangle$, with $x = (x^1, \dots, x^d)$ in d dimensions. In the absence of any potential ($H_{\text{cl}} = 0$), the Schrödinger equation in imaginary time is just describing the free diffusion of this particle:

$$\frac{\partial \psi(x, \tau)}{\partial \tau} = \frac{1}{2m} \nabla^2 \psi(x, \tau) \quad (3.15)$$

for a particle of mass m . If one considers the particle initially localized in space, $\psi(x, 0) = \delta(x - x_0)$, the corresponding solution of Eq. (3.15) (its Green function) is the well-known heat kernel

$$G_{\text{qu}}(x, x_0, \tau) = \left(\frac{m}{2\pi\tau} \right)^{d/2} e^{-\frac{m}{2\tau}(x-x_0)^2}, \quad (3.16)$$

as can be easily shown by solving the equation in momentum space. As usual, we can use the Green function as a real-space propagator to generate the solution of Eq. (3.15) for any initial condition:

$$\psi(x, \tau) = \int dx' G_{\text{qu}}(x, x', \tau) \psi(x', 0). \quad (3.17)$$

On the other hand, the action of the classical part is diagonal in real space, $H_{\text{cl}} = V(x)$, so it simply acts as a multiplication operator, $(e^{-\tau H_{\text{cl}}} \psi)(x, \tau) = e^{-\tau V(x)} \psi(x, \tau)$. In other words, we have the propagator

$$G_{\text{cl}}(x, x_0, \tau) = \delta(x - x_0) e^{-\tau V(x)}. \quad (3.18)$$

The factorization property Eq. (3.14) implies that as long as $\delta\tau = \tau/K$ is small, the full propagator corresponding to $H = H_{\text{cl}} + H_{\text{qu}}$ is well approximated (*i.e.* to order $O(\delta\tau^2)$) by the concatenation

$$G(x, x', \delta\tau) \approx G_{\text{qu}}(x, x', \delta\tau) G_{\text{cl}}(x, x', \delta\tau). \quad (3.19)$$

In order to compute $\psi(x, \tau) = \int dx' G(x, x', \tau) \psi(x, 0)$, the DMC algorithm implements a classical stochastic dynamics on a population of “walkers”, whose distribution at time τ is described by $\psi(x, \tau)$. This basically consists in switching from a Fokker–Planck description of the system (*i.e.* in terms of the *deterministic* evolution of a probability distribution) to a Langevin description (in terms of the *stochastic* evolution of a particle).

More precisely, we start the simulation by initializing N_w walkers, *i.e.* points in configuration space, according to a certain initial profile $\psi(x, 0)$. Typical choices for the population size range within $N_w \approx 10^2 - 10^4$, with larger numbers resulting in a smaller systematic errors. Now consider an individual walker. In the single-particle setting, we can regard this as a classical particle in real space. The free diffusion described by the quantum part of the propagator can be implemented by randomly displacing the walker from its original position x_0 such that the new position x is distributed according to $G_{\text{qu}}(x, x_0, \delta\tau)$ (seen as a probability distribution for the variable x). This corresponds to a step in a random spatial direction whose length is a Gaussian random variable with variance $\delta\tau$.

The classical part of the propagator ought to operate a reweighting procedure according to the local value of the potential $V(x)$. In other words, the Langevin dynamics should award walkers for staying in low-potential regions, and penalize them for wandering off to higher grounds. In principle, one could simply try to kill the walker with a probability determined by the classical weight $w(x) \propto e^{-\delta\tau V(x)}$. However, this would decimate the walker population far too quickly (indeed exponentially so), making the simulation impossible. This effect is the manifestation, at the walker level, of our previous observation that the

imaginary time evolver is not unitary. There is a standard [248] procedure to obviate this difficulty. We can introduce a *branching* rule whereby walkers can be duplicated as well as deleted, in such a way that the average replication rate reproduces the correct weight.

A way to achieve this is the following: in between diffusion steps, the weight $w(x)$ for the walker is computed. Then, a uniform random number $r \in [0, 1]$ is drawn. Finally, the walker is replaced by $\lfloor w(x) + r \rfloor$ copies of itself in the same spatial location.³ One also has to properly select the normalization for the weight $w(x)$. This is equivalent to choosing a “reference energy” E_T so that the weight has the form

$$w(x) = e^{-\delta\tau(V(x)-E_T)}. \quad (3.20)$$

For a proper execution, E_T should have the effect of fixing the average reproduction rate to about 1, so as to avoid both an exponential decay and an exponential growth of the walker population. In practice, in order to keep the number of walker under control, E_T must be adjusted after every step. One preemptively fixes a target number of walkers $N_{w,\text{target}}$ and every instantaneous deviation from it is negatively fed back to E_T , *e.g.* via the update rule

$$E_T \mapsto E_T + \alpha \log\left(\frac{N_{w,\text{target}}}{N_w(\tau)}\right) \quad (3.21)$$

for some $\alpha > 0$. In alternative, it may be preferable to allow for small drifts from the target by using the instantaneous deviation, namely replacing $N_{w,\text{target}}$ with $N_w(\tau - \delta\tau)$ in Eq. (3.21). We can see from Eq. (3.12) that the ideal value for E_T , canceling out any exponential behavior of the state’s norm, coincides with the ground energy E_0 . Therefore, the stabilized value of the regulator can be used as an estimate for the ground energy.

Overall, the master equation for the walker distribution reads like a diffusion equation with an additional source term:

$$\frac{\partial\psi(x, \tau)}{\partial\tau} = \frac{1}{2m}\nabla^2\psi(x, \tau) - [V(x) - E_T]\psi(x, \tau). \quad (3.22)$$

The methods described so far can be applied to many-body systems without much modification. The real-space basis $\{|x\rangle\}_x$ will be traded for a computational basis $\{|z\rangle\}_z$, *e.g.* the standard σ^z basis for a spin system, with the potential part acting diagonally on the many-body state. Walkers will correspond to computational basis states, moving in an abstract (*e.g.* the Boolean hypercube for spins- $\frac{1}{2}$) as opposed to real space. The only nontrivial adjustment concerns the diffusion part of the propagator, whose corresponding transfer matrix needs to be computed for the specific model as in Eq. (3.6). For the Ising model, for example, one obtains the following expression:

$$G_{\text{qu}}(z, z', \delta\tau) = \frac{\sinh(\delta\tau\Gamma)^r \cosh(\delta\tau\Gamma)^{n-r}}{\exp(n\delta\tau\Gamma)}, \quad (3.23)$$

³This is equivalent to allowing a walker to survive with probability $w(x)$, plus creating a new one with probability $w(x) - 1$ should the weight exceed 1, plus another one with probability $w(x) - 2$ should it exceed 2, *etc.* [107]. It is usually wise to restrict the maximum number of allowed copies per generation.

where $r = |z - z'|_{\text{H}}$ is the Hamming distance (minimum number of spin flips) between the configurations, and Γ the transverse field intensity. This is properly normalized as $\sum_{z=0}^{D-1} G_{\text{qu}}(z, z', \delta\tau) = 1$ for all $z' \in \{0, 1\}^n$.⁴

As alluded to previously, the goal of DMC is to approximate the ground-state wave function $\psi_{\text{gs}}(z) = \langle z | E_0 \rangle$. Using the formal analogy between the imaginary-time Schrödinger equation and a diffusion equation, the wave function $\psi(z, \tau)$ of a system is interpreted as the probability distribution of a population of walkers evolving under a specific Langevin dynamics. However, this analogy only makes sense insofar as the wave function can be treated as a probability distribution. This cannot be done in general, as $\psi(z, \tau)$ need not be nonnegative (we can however always assume that it is real).

That said, as we mentioned toward the end of the previous section, the ground state of *stoquastic* Hamiltonians (cf. Eq. (3.11)) can always be taken to be nonnegative,

$$\psi_{\text{gs}}(z) \geq 0 \quad \forall z, \quad (3.24)$$

and hence treated as a probability (rather than amplitude) distribution.

Importance sampling

Unless the potential $V(z)$ is particularly well-behaved, the weight factor $w(z) = e^{-\tau(V(z) - E_{\text{T}})}$ will vary with z regardless of the regulator E_{T} . This means that, although one is free to fix the average replication rate, the distribution of $\lfloor w(z) + r \rfloor$ will in general cause the walker population to fluctuate considerably.

A ubiquitous approach to sampling irregular functions through Monte Carlo methods is *importance sampling*. It consists in using a function which approximates the desired distribution as a “guide” for the sampling process. In the case of DMC, we introduce a *guiding wave function* (GWF) $\Psi_{\text{G}}(z)$ that we expect to decently approximate the ground state, $\Psi_{\text{G}}(z) \approx \psi_{\text{gs}}(z)$. We then shift our aim toward sampling the *product*

$$\rho(z, \tau) = \Psi_{\text{G}}(z)\psi(z, \tau). \quad (3.25)$$

Let us go back to the single-particle case for the sake of simplicity. The addition of a GWF in Eq. (3.22) leads to the appearance of an extra drift term in the master equation (which we write in one dimension to keep the notation lighter):

$$\frac{\partial \rho(x, \tau)}{\partial \tau} = \frac{1}{2m} \frac{\partial^2 \rho(x, \tau)}{\partial x^2} + \frac{1}{m} \frac{\partial}{\partial x} [\varphi'_{\text{G}}(x)\rho(x, \tau)] - [E_{\text{L}}(x) - E_{\text{T}}]\rho(x, \tau), \quad (3.26)$$

with the *local energy* $E_{\text{L}}(x)$ defined by

$$E_{\text{L}}(x) = V(x) - \frac{1}{2m\Psi_{\text{G}}(x)} \frac{d^2\Psi_{\text{G}}(x)}{dx^2}, \quad (3.27)$$

⁴To see this, call $p = \sinh(\delta\tau\Gamma)/e^{\delta\tau\Gamma}$ and notice that $G_{\text{qu}}(z, z', \delta\tau) = p^r(1-p)^{n-r}$ with $r = |z - z'|_{\text{H}}$. The sum over z can be written as a sum over the Hamming distance r times a binomial coefficient $\binom{n}{r}$, which makes the general terms into a binomial distribution.

the effective potential by

$$\varphi_G(x) = -\log|\Psi_G(x)| \quad (3.28)$$

and we used the notation $\varphi'_G = \frac{d\varphi_G}{dx}$. Notice how Eq. (3.26) reduces to Eq. (3.22) when $\Psi_G(x) = 1$ (no importance sampling).

As is clear from Eq. (3.26), the effect of a nontrivial GWF is twofold: first, it adds a deterministic drift $\delta x = \delta\tau \varphi'_G(x)$ on top of the diffusion process, which pushes walkers toward regions of larger $|\Psi_G(x)|$, thus facilitating the convergence process. Secondly, it reshapes the local potential felt by the walkers from $V(x)$ into $E_L(x)$. The local energy is a constant function whenever the GWF is an exact eigenfunction of the original Hamiltonian (indeed it is the eigenvalue corresponding to Ψ_G). In this case, there is no reweighting during the evolution, and the walker population is constant at all times. In a realistic case, the GWF will not be an eigenfunction, but it should be close (in some functional norm) to the exact ground state. This means that $E_L(x)$ will be “almost a constant”, *i.e.* much better behaved than the naked potential $V(x)$, resulting in much smaller fluctuations in the walker population and a proportionally stabler algorithm.

Therefore, in order for DMC to be effective one needs to be able to produce a satisfactory ansatz for the ground state. This is typically obtained variationally, *i.e.* by writing a candidate wave function with some free parameters and then minimizing its energy with respect to them. We are going to see a couple of simple examples in the course of the next section, but generally speaking, the problem of finding good variational approximants to many-body ground states is extremely complex, and a major subject of research in condensed matter, nuclear physics, quantum chemistry, and more.

3.3 Simulating tunneling events on classical machines

In this section, we try to apply quantum Monte Carlo methods, specifically DMC, to problems which we expect to be efficiently tackled by QA, on account of their energy landscapes containing tunneling paths. We begin by contextualizing and motivating this investigation in Section 3.3.1, before proceeding to the study of a single-particle model on the real line — the quartic double well model — which will serve as a testing ground for our methods (Section 3.3.2). Despite its conceptual simplicity, this model can be seen as a mean-field approximation to a large class of many-body systems characterized by a single relevant parameter which can switch between two different equilibrium values when activated over, or tunneling through, a tall barrier. Examples include the total magnetization in 1D spin models and reaction coordinates in chemistry. Moreover, this very simplicity will allow us to present a semiclassical theory explaining the observed phenomenology (Section 3.3.3). Afterwards, we will tackle two many-body spin systems on a lattice, namely the familiar Ising model in a transverse field (Section 3.3.4) and the so-called “shamrock model” (Section 3.3.5). Specific attention will be given in all cases to the role of importance sampling in the performance of the algorithm. We will comment on

the systematic error introduced by using a finite walker population, and the complexity requirements that stem therefrom, in Section 3.3.6. Finally, we will present our conclusions in Section 3.3.7.

The work presented in Sections 3.3.2 and 3.3.3 was performed in collaboration with Giuliano Giudici, and that in Sections 3.3.4 through 3.3.6 was carried out by Estelle M. Inack. The project as a whole was conceived and supervised by Sebastiano Pilati.

3.3.1 Previous results and motivation

As we saw in Section 1.3.2, quantum annealers are special purpose adiabatic quantum computers designed to solve complex optimization problems. Compared to alternative classical optimization algorithms, chiefly simulated annealing, they can additionally exploit quantum tunneling to cross energy barriers and reach lower energy solutions. Their dominant bottlenecks are the small energy gaps associated to avoided level crossings, which typically occur in disordered systems when two well-separated competing states are connected by a tunneling process. This scenario frequently happens in the glassy phases that characterize typical hard optimization problems.

The task of classically simulating the real-time dynamics of quantum annealers, *e.g.* to identify classes of problems where they might outperform classical optimization methods, is only possible for relatively small systems (say, around 30 spins). Quantum Monte Carlo (QMC) simulations, making use of algorithms such as those described in Section 3.2, have emerged as a useful alternative tool to simulate the quantum annealers' behavior in configurations where the sign problem does not occur [12, 49, 104, 120, 170, 175, 242]. This is the case, for instance, of the devices currently commercialized by D-Wave systems (see *e.g.* Refs. [48, 49, 137, 157, 218]). In particular, PIMC [65] and projective QMC (PQMC) [107] algorithms have been adopted, beside other techniques such as the stochastic series expansion algorithm [160, 161].

Tunneling events similar to those characterizing the quantum annealers' dynamics also occur during QMC simulations [50, 80]. In various problems characterized by a double-well energy landscape, the tunneling rate of nonzero-temperature PIMC simulations was found to scale with the system size, or with the height of the energy barrier, as the square of the first energy gap [55, 133, 174]. This is the same scaling predicted by the theory of incoherent quantum tunneling [262], and it is also the scaling of the inverse of the annealing time required by a coherent quantum annealer to avoid diabatic transitions [100].

Refs. [133, 136, 174] explained these results using a semiclassical theory of instantons in PIMC simulations. In the case of PQMC, we find in Sections 3.3.2–3.3.5 the tunneling rate to scale linearly with the gap, providing a quadratic speedup compared to the expected behavior of a quantum annealer.⁵ Furthermore, the PQMC algorithms displays

⁵A linear scaling was identified also for PIMC simulations performed with open boundary conditions in the inverse-temperature direction [133, 174]. However, significant deviations have later been discussed [136]. Furthermore, the computational cost of zero-temperature simulations based on open-boundary PIMC algorithms

this speedup even in the so-called shamrock model (Section 3.3.5), where frustrated interactions cause an exponential slowdown of the finite-temperature PIMC dynamics [22]. These findings suggest that PQMC simulations constitute a relevant benchmark for physical quantum annealers and a competitive quantum-inspired classical optimization algorithm [71, 131, 135]. In fact, they have recently been employed to obtain better solutions in optimization problems relevant for medical research, specifically, for pulse-sequence optimization in magnetic resonance fingerprinting [156]. This further highlights the importance of exhaustively characterizing their tunneling dynamics.

The tunneling-time studies mentioned above have considered PQMC algorithms implemented without a guiding wave function (GWF). The GWF, usually a variational ansatz that closely approximates the ground state, guides the PQMC simulation towards the relevant regions of the configuration space. This improves the algorithm's accuracy and efficiency [107]. In fact, without a sufficiently accurate GWF, the computational cost of PQMC simulations increases exponentially with the system size [52, 193]. In principle, one might expect the GWF to significantly impact the tunneling dynamics, since it alters both the sampling algorithm and the probability distribution sampled at equilibrium. In this section, we analyze whether the GWF does indeed affect the tunneling time in PQMC simulations, and if it does, to what extent. As test beds, we consider a one-dimensional continuous-space Hamiltonian describing a quantum particle in a double well, the ferromagnetic quantum Ising chain, and the shamrock model. Notice that, in the ferromagnetic phase, also the Ising-type models can be described by an effective double-well potential, with the two polarized states with opposite magnetizations representing the two competing potential minima. We consider different kinds of GWFs, including a Boltzmann-type ansatz that mimics the equilibrium distribution of a classical statistical ensemble and, for the continuous-space model, the numerically exact representation of the ground state. For the quantum Ising chain, we also consider an ansatz that mimics the structure of a generative artificial neural network [62], specifically an unrestricted Boltzmann machine [132] or, in a different jargon, a shadow wave function [214, 257]. Remarkably, for all GWFs we consider, we find the same linear scaling (to leading exponential order) of the tunneling rate with the gap as previously found in PQMC simulation performed without GWF. The choice of the GWF only affects the prefactor. We also provide a semiclassical theory based on the Wentzel–Kramers–Brillouin (WKB) approximation, valid for PQMC simulations of the double well with exact GWF, that explains the observed linear scaling.

has not been analyzed in detail.




Figure 3.1: Profile of the quartic double-well potential Eq. (3.30) (dotted blue line), shifted for better comparison with the wave functions of the ground state $\Psi_0(x)$ and of the first excited state $\Psi_1(x)$ (orange and green solid lines, respectively).

3.3.2 DMC tunneling in a quartic double-well potential

In this section, we consider a quantum particle in one spatial dimension, described by the following continuous-space Hamiltonian:

$$\hat{H} = -\frac{1}{2} \frac{d^2}{dx^2} + V(x), \quad (3.29)$$

with the quartic double-well potential

$$V(x) = \frac{x^4}{g} - x^2. \quad (3.30)$$

The profile of $V(x)$ is visualized in Fig. 3.1, together with the corresponding ground-state wave function $\Psi_0(x)$ and the first excited state $\Psi_1(x)$. Units are chosen so that $\hbar/m = q = 1$, where \hbar is the reduced Planck constant, m is the particle mass and $q = \frac{1}{2}\sqrt{V''(x_{\ell,r})}$ fixes the curvature at the bottom of the well. Here $x_{\ell,r} = \mp\sqrt{g/2}$ are the minimum points of $V(x)$. The dimensionless parameter g controls the height of the barrier separating the two wells, $\Delta V = V(0) - V(x_\ell) \propto g$, and the distance between the minimum points, $x_r - x_\ell \propto \sqrt{g}$. For large g , the barrier height increases relative to the wells' separation and the energy spectrum becomes doubly degenerate, corresponding to the two wells being asymptotically independent. For large but finite g , the two wells are connected by tunneling processes. These processes lift the degeneracies, leading to small tunneling gaps

between the energy levels. An approximate expression for the first energy gap $\Delta = E_1 - E_0$ between the ground state energy E_0 and the first excited level E_1 , valid in the large- g regime, can be determined via the WKB semiclassical theory, which gives [110]

$$\Delta \sim 8\sqrt{\frac{g}{\pi}} \exp\left(-\frac{2g}{3}\right). \quad (3.31)$$

In an isolated double well, a quantum particle initially prepared in one of the two states

$$\psi_{\ell,r}(x) = \frac{\Psi_0(x) \pm \Psi_1(x)}{\sqrt{2}}, \quad (3.32)$$

which are localized in the left and right wells, respectively, performs coherent Rabi oscillations between the two states at a rate proportional to Δ . As a reference, it is worth mentioning that, instead, a quantum particle coupled to a thermal bath and subjected to a double-well potential would undergo *incoherent* tunneling at a rate proportional to Δ^2 [262].

Tunneling time in DMC simulations

Our goal is to analyze the relation between the *DMC tunneling time* ξ , defined as the imaginary time required by the walkers to leak from one well to the other, and the physical tunneling time of the real-time dynamics. More precisely, we are interested in the scaling relation between ξ and the inverse energy gap Δ^{-1} . We measure ξ with a protocol analogous to the one adopted in Refs. [133, 174] for PIMC simulations. All walkers are initially set at the bottom of the left well $x = x_\ell$. The DMC simulation is run until a percentage p of the instantaneous walker population overcomes a given position threshold in the right well $x_{\text{thr}} \geq 0$. The final imaginary time is recorded, and the process is repeated approximately 300 times to accumulate statistics. The fluctuations of the final imaginary times turn out to be approximately normally distributed, and we take the average and its standard deviation as the definition of ξ and of its error bar, respectively. In the simulations reported here, the threshold is set at $x_{\text{thr}} = x_r/2$ and the walker percentage at $p = 25\%$. A careful analysis shows that the asymptotic scaling of ξ is independent of this specific choice of p and x_{thr} up to a constant prefactor. Furthermore, the chosen target walker population $N_w \approx 10^4$ is large enough, and the chosen time step small enough (e.g. $\delta\tau = 0.007$ for $1/\Delta > 70$), to eliminate any significant systematic error on ξ .

First, the DMC algorithm is run without a guiding wave function (GWF). We measure the tunneling time ξ for different barrier heights, tuned by varying the parameter g . In Fig. 3.2, ξ is plotted as a function of the inverse energy gap Δ^{-1} , which we compute for the different g values using a standard finite-difference method [68]. The discretization is




Figure 3.2: DMC tunneling time ξ for the quartic double-well potential (3.30) as a function of the inverse energy gap Δ^{-1} . Three different DMC protocols are shown: the simple DMC algorithm without a GWF (red circles), the DMC algorithm guided by a Boltzmann ansatz (green empty squares), and the one guided by the numerically-exact representation of the ground state $\Psi_0(x)$ (blue empty diamonds). The dashed line represents the scaling $\xi \propto \Delta^{-1}$. Here and in all plots, if not visible, the error bars are smaller than the symbol size.

fine enough to ensure that there is no sizable finite-precision effect. In the large- g regime, corresponding to large Δ^{-1} , the tunneling times approach the scaling law

$$\xi \propto \Delta^{-1}. \quad (3.33)$$

We are going to find that the same relation holds for the Ising-type models considered in Section 3.3.4.

Next, we run DMC simulations with a GWF. First, we consider as GWF the approximate “Boltzmann-type” ansatz⁶

$$\Psi_G(x) = \exp[-\beta V(x)]. \quad (3.34)$$

The fictitious inverse temperature β is fixed by minimizing the variational energy estimate

$$E_G(\beta) = \frac{\langle \Psi_G | H | \Psi_G \rangle}{\langle \Psi_G | \Psi_G \rangle}. \quad (3.35)$$

⁶Notice that $\Psi_G(x)$ does not need to be normalized to one, as it only enters the master equation (3.26) through ratios.

Second, we consider a numerical representation of the exact ground-state wave function, *i.e.*, we set $\Psi_G(x) = \Psi_0(x)$. The ground-state wave function $\Psi_0(x)$ is obtained via the same finite-difference technique used for the gaps. Notice that this ansatz represents the optimal GWF for equilibrium simulations.⁷

In Fig. 3.2, the tunneling times obtained with these two GWFs are compared with the results obtained without any. Remarkably, for large g (tall barrier), the *very same* linear relation between ξ and Δ^{-1} is approached. By fitting the three datasets in the large- g regime with the function $\xi(\Delta) = \alpha\Delta^{-b}$, where α and b are the fitting parameters, we obtain the values reported in Table 3.1. In all three cases, the exponent b is consistent with the linear relation corresponding to $b = 1$. The choice of GWF only affects the prefactor α , though for this model the variations are small enough to be masked by statistical uncertainties. These findings indicate that the GWF does not affect the leading scaling relation between tunneling time and inverse energy gap. This is a surprising results, given that introducing the GWF affects both the sampling algorithm and the equilibrium probability distribution of the DMC simulation. A rough explanation can be conjectured by considering the competition between the two effects originating from the introduction of the GWF. The first is due to the deterministic drift, which pushes walkers away from the potential barrier, inhibiting inter-well crossings. The second is the smoothing out of the weight reduction that occurs when walkers encounter a bump in the potential; this effect reduces the probability of those walkers being eliminated from the population, enhancing tunneling. Our numerical results indicate that these two effects tend to compensate exactly. A more formal explanation of the relation $\xi \sim \Delta^{-1}$ is proposed in the next subsection.

The double-well potential (3.30) is characterized by a specific functional relation between the barrier's height and width, specifically

$$\frac{\text{height}}{\text{width}} = \frac{V(0) - V(x_r)}{x_r - x_\ell} = \sqrt{\frac{g}{32}}. \quad (3.36)$$

However, we expect our findings to hold for more generic choices of double-well potentials as well. In order to verify this, we introduce an adjustable parameter that allows us to vary the width of the barrier independently of its height. Similarly to Ref. [174], we consider the potential

$$U(x) = \frac{(|x| - x_0)_+^4}{g} - (|x| - x_0)_+^2, \quad (3.37)$$

where $f(x)_+ \equiv \max\{0, f(x)\}$ and $x_0 \geq 0$. This potential features a *plateau* of width $2x_0$ around the origin, and reduces to $V(x)$ for $x_0 = 0$. In our study, the barrier height is kept

⁷One could argue that using the ground state as a GWF for DMC makes as much sense as sketching out a suspect's identikit by tracing over a high-resolution photograph. While it is true that in realistic settings one never does have access to the exact ground state — nor would he need DMC if he did —, now we are only interested in understanding the tunneling dynamics. By benchmarking against the best possible GWF, we can upper bound the performance gain from all the others.

DMC	α	b
No GWF	109(11)	0.99(2)
Boltzmann	98(16)	1.01(3)
Exact GWF	112(8)	0.99(1)

Table 3.1: Fitting parameters α and b , describing the small-gap behavior of the DMC tunneling time ξ in the double-well potential (3.30) according to the fitting function $\xi(\Delta) = \alpha\Delta^{-b}$. For each protocol, the five rightmost data points shown in Fig. 3.2 are included in the fit.

constant by fixing g to some value (we use $g = 8$), while x_0 is increased in the interval $x_0 \in [0, 2]$. This has the effect of reducing the gap Δ as well as the tunneling rate ξ^{-1} . The GWF chosen for this study is the numerically exact ground-state wave function. Once again, a linear relation is found between ξ and Δ^{-1} , and fitting the dataset with the power law $\xi(\Delta) = \alpha\Delta^{-b}$ in the small-gap regime yields the values $\alpha = 23(1)$ and $b = 0.993(7)$ for the parameters, confirming the generality of relation (3.33).

3.3.3 Semiclassical theory of the DMC tunneling dynamics

We now present a semiclassical theory to explain and generalize our findings. As mentioned above, the DMC algorithm with GWF is described by Eq. (3.26), containing both the usual drift and diffusion terms as well as an additional norm-nonpreserving term from the branching process. When the GWF coincides with the exact ground-state wave function, $\Psi_G(x) = \Psi_0(x)$, the branching rate becomes a constant and can be set to zero by a suitable choice of E_T . This results in the standard Fokker–Planck equation

$$\frac{\partial \rho(x, t)}{\partial t} = \frac{1}{2} \frac{\partial^2 \rho(x, t)}{\partial x^2} + \frac{\partial}{\partial x} [\varphi'_G(x) \rho(x, t)]. \quad (3.38)$$

This equation describes the stochastic dynamics of a classical particle with distribution $\rho(x, t)$ subject to the effective potential $\varphi_G(x)$ defined in Eq. (3.28) (with $\Psi_G = \Psi_0$). The tunneling time corresponding to this dynamics can then be identified (to exponential accuracy) with the activation time needed for this classical ensemble to overcome the effective barrier $\Delta\varphi_G \equiv \varphi_G(0) - \varphi_G(x_{\min})$, with initial conditions $\rho(x, 0) = \delta(x - x_\ell)$ (the normalization can be set to one since the norm of $\rho(x, t)$ is conserved). Here, x_{\min} indicates the (left) minimum point of $\varphi_G(x)$, in general different from x_ℓ , while $x = 0$ is its local maximum point, as follows from the parity of $\Psi_0(x)$. The computation of the classical activation over a potential barrier is known as Kramers problem [140, 151]. The corresponding activation time reads

$$\tau_{\text{act}} = \frac{2\pi}{\sqrt{\varphi''_G(x_{\min}) |\varphi''_G(0)|}} \exp(2\Delta\varphi_G). \quad (3.39)$$

The large- g scaling of τ_{act} is dominated by the exponential function in Eq. (3.39), so we focus on that term alone. Substituting Eq. (3.28) into Eq. (3.39), one obtains

$$\tau_{\text{act}} \sim \left(\frac{\Psi_0(0)}{\Psi_0(x_{\min})} \right)^{-2} \sim \Psi_0(0)^{-2}. \quad (3.40)$$

Here, we used the fact that $\Psi_0(x_{\min})$ cannot be exponentially small in g if the ground-state wave function is to be normalized to unity (notice that, by definition, $\Psi_0(x)$ achieves its global maximum at x_{\min} , and its amplitude is exponentially small outside of a region of width $O(\sqrt{g})$). The value of $\Psi_0(0)$ can be estimated using WKB theory, which gives $\Psi_0(0) \sim \exp(-g/3)$ [191]. Substituting this into Eq. (3.40), and combining with the WKB estimate of the gap, Eq. (3.31), we finally get

$$\xi \sim \tau_{\text{act}} \sim \frac{1}{\Delta}, \quad (3.41)$$

where the symbol \sim denotes asymptotic equal scaling for $g \rightarrow \infty$, up to subexponential corrections.

The above argument can actually be generalized to double-well-type potentials other than the quartic double well defined in Eq. (3.30). In fact, Eq. (3.40) can be derived from Eq. (3.38) under quite general assumptions on $V(x)$, as long as a sufficiently accurate approximant of the ground-state wave function is available to use as a GWF, such that the effect of branching may be neglected. Moreover, for a generic double-well potential the ground gap in the large-barrier limit can be expressed as [110, 155]:

$$\Delta \propto \psi_r(0)\psi_r'(0), \quad (3.42)$$

where $\psi_r(x)$ is the ground-state wave function of a *single* well, as defined in Eq. (3.32). The above relation holds in the only hypothesis that $\psi_r(x)$ is asymptotically localized on $x > 0$, such that the condition

$$\int_{-\infty}^0 [\psi_r(x)\psi_r(-x) - \psi_r(x)^2] dx \ll 1 \quad (3.43)$$

is fulfilled at large g . We can now resort to WKB theory, which gives the following expression for $\psi_r(x)$ (to exponential accuracy):

$$\psi_r(x) = \exp\left(-\int_x^a dx' \sqrt{2(V(x') - E_0)}\right), \quad (3.44)$$

where E_0 is the ground-state energy and a is the (positive) classical turning point, defined by $V(a) = E_0$. This implies that, up to subexponential corrections, $\Delta \propto \psi_r(0)^2 \propto \Psi_0(0)^2$. Upon substitution into Eq. (3.40), this leads to the scaling relation (3.41). Therefore, in the assumption that the WKB approximation works well in a neighborhood of the origin (e.g. assuming that the turning points do not approach 0 in the infinite-barrier limit), we see that the DMC tunneling time generically exhibits a Δ^{-1} scaling regardless of the specific form of the potential.

3.3.4 The ferromagnetic quantum Ising chain

We now consider interacting spin models on a lattice. Our DMC simulations for discrete-basis models are based on the continuous-time Green function Monte Carlo algorithm [237], exhaustively described in Ref. [39].

In this section we present the results for the one-dimensional ferromagnetic transverse-field Ising model (TFIM) defined by the following Hamiltonian:

$$H = -J \sum_{i=0}^{L-1} \sigma_i^z \sigma_{i+1}^z - \Gamma \sum_{i=1}^L \sigma_i^x, \quad (3.45)$$

with a ferromagnetic coupling $J > 0$. Periodic boundary conditions are considered, $\sigma_L^z = \sigma_0^z$.

At zero temperature this model undergoes a quantum phase transition from a paramagnetic phase for $\Gamma > J$ to a ferromagnetic phase for $\Gamma < J$. In the $\Gamma \rightarrow 0$ limit quantum fluctuations are suppressed and one has two degenerate (classical) states with all spins up $|\uparrow\uparrow \dots \uparrow\rangle$ or all spins down $|\downarrow\downarrow \dots \downarrow\rangle$. In order to flip from one state to the other, the system must overcome a free energy barrier separating the two minima, with the magnetization playing the role of a one-dimensional reaction coordinate which parametrizes a symmetric double-well profile of the kind described in the previous sections. For small $\Gamma > 0$, in the thermodynamic limit there are still two degenerate ground states with opposite magnetizations, but in a finite chain the degeneracy is lifted by an exponentially small (in the system size) energy gap due to the quantum tunneling coupling the two states. This closing-gap scenario resembles the Landau–Zener avoided level crossings one typically encounters in adiabatic quantum optimization. There, the small gaps are associated to tunneling processes between competing solutions. In order to avoid diabatic transitions to the first excited state, the total annealing time has to scale as $1/\Delta^2$. These small gaps represent the bottleneck of adiabatic quantum computing, since for hard optimization problems they often close exponentially fast with the problem size [33, 99, 146, 158, 267].

We define the DMC quantum tunneling time ξ analogously to the one-dimensional case, except we now stop the clock as soon as $p = 25\%$ of the walkers (which are up/down configurations of the Ising chain), initialized with fully positive magnetization, come to have a negative magnetization (majority of spins pointing down), meaning that they have crossed the free energy barrier. This definition is analogous to the one employed in Refs. [22, 133, 174] in the case of PIMC simulations, where a certain percentage of imaginary-time slices, instead of walkers, is considered. The simulation is repeated approximately 250 times for larger systems and small Γ and approximately 2 500 for smaller systems and larger values of Γ . As in the double-well case, we then take the average value to define ξ and its standard deviation to define the error bar.

The DMC tunneling times for the ferromagnetic Ising chain are shown in Fig. 3.3, as a function of the number of spins L and for different values of Γ . For large L the data display

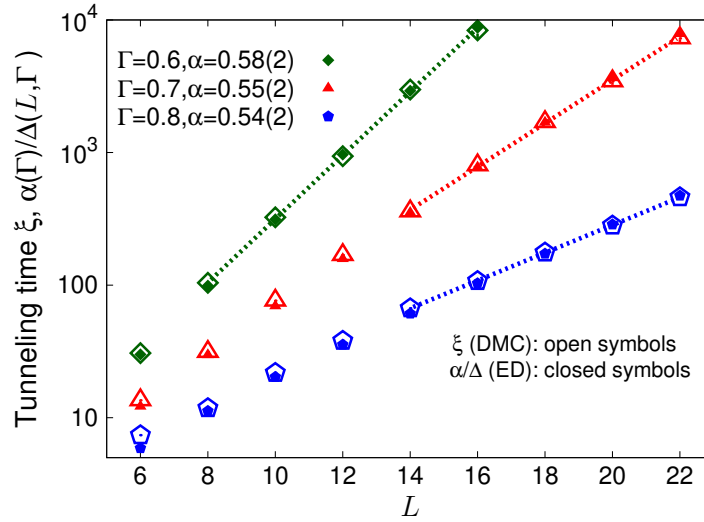


Figure 3.3: DMC tunneling time ξ for the ferromagnetic Ising chain (open symbols) as a function of the number of spins L , for different values of the transverse field Γ with $J = 1$. The closed symbols represent the inverse gap values $1/\Delta$ obtained with exact diagonalization and rescaled by a parameter $\alpha(\Gamma) \sim 1$ (see legend). The thin dashed curves represent exponential fits for the tunneling time ξ in the large- L regime. Here and in the other graphs the error bars are smaller than the symbol size if not visible.

an exponential growth, quite similar to the dependence of the inverse gap $1/\Delta$, which we obtain via exact diagonalization of the Hamiltonian matrix. In fact, by multiplying the inverse gap $1/\Delta$ by an appropriate numerical prefactor α we obtain precise matching between the two datasets. The coefficient α turns out to be a number of order 1. We also consider different definitions of DMC tunneling time, using percentages of walkers that have to cross the barrier between 10% and 25%, obtaining again results which follow the $1/\Delta$ scaling but with a slightly different value of the prefactor α .

The inverse-gap scaling displayed by the DMC tunneling times is similar to the result found in Ref. [133] using modified PIMC simulations performed using open boundary conditions in imaginary time. This is not surprising, since such modified PIMC method had been originally introduced as a computational tool to study ground-state properties [72,222]. However, it is usually employed in combination with guiding wave functions that accurately describe the ground state, so that the convergence to the zero-temperature limit as a function of the total path length is quite rapid. How this algorithm converges to the ground state in the absence of a GWF has not been analyzed in detail yet. It is also worth stressing that in the PIMC formalism the tunneling time is defined by counting the number of Monte Carlo sweeps (a sweep corresponds to one Monte Carlo step per spin) and, therefore, it does not bear the significance of imaginary time as in the DMC method employed in this article. In Ref. [133], finite-temperature PIMC simulations (with

periodic boundary conditions) have also been performed, finding that the PIMC tunneling times scale as $1/\Delta^2$. This behavior was found in ferromagnetic Ising models (which are characterized by a one-dimensional reaction coordinate) and it was later confirmed in one-dimensional and two-dimensional continuous-space models as well [174], showing that it persists even when the reaction coordinate is multidimensional.

Considered together, the above findings suggest that QMC algorithms are either as efficient as (in the case of finite-temperature PIMC) or quadratically faster than (in the case of PIMC with open boundary conditions in imaginary time, and DMC) quantum annealers in tunneling through free energy barriers and therefore, if one assumes that incoherent quantum tunneling is the empowering resource of QAs, also in solving optimization problems.

TFIM with importance sampling

In this section we consider two types of GWF. The first one is the Boltzmann ansatz:

$$\Psi_G(z) = \exp[-\beta E_{\text{cl}}(z)], \quad (3.46)$$

which is modeled on the Boltzmann distribution of a classical Ising system with Hamiltonian function $E_{\text{cl}}(z) = \langle z | H_{\text{cl}} | z \rangle$, where $H_{\text{cl}} = -J \sum_i \sigma_i^z \sigma_{i+1}^z$. The fictitious temperature β is fixed by minimizing the variational energy $\frac{\langle \Psi_G | H | \Psi_G \rangle}{\langle \Psi_G | \Psi_G \rangle}$ using the stochastic gradient descent method [237].

The second GWF is based on a stochastic generative neural network, specifically an unrestricted Boltzmann machine (uRBM) [132]; it is defined as

$$\Psi_G(z) = \sum_h \phi(z, h) \quad (3.47)$$

where

$$\phi(z, h) = \exp \left[\sum_{i=0}^{L-1} (K_1 z_i z_{i+1} + K_2 h_i h_{i+1} + K_3 z_i h_i) \right]. \quad (3.48)$$

The wave function amplitude $\Psi_G(z)$ in each physical, or *visible*, spin configuration $z = (z_0, \dots, z_{L-1})$ is obtained by integrating over all configurations of the L *hidden units* $h = (h_0, \dots, h_{L-1})$, which take the values $h_i \in \{\pm 1\}$ for $i \in \{0, \dots, L-1\}$. Periodic boundary conditions are considered both in the visible and in the hidden layers, $z_L = z_0$ and $h_L = h_0$. The three coupling constants K_1 , K_2 , and K_3 fix the interaction strengths between nearest-neighbor visible–visible and hidden–hidden pairs, and between visible–hidden pairs with the same index, respectively. We determine them via variational optimization using the stochastic reconfiguration method.

The analysis reported in Ref. [132] indicated that optimized uRBM GWFs are sufficiently accurate to considerably reduce the computational cost of DMC simulations of the

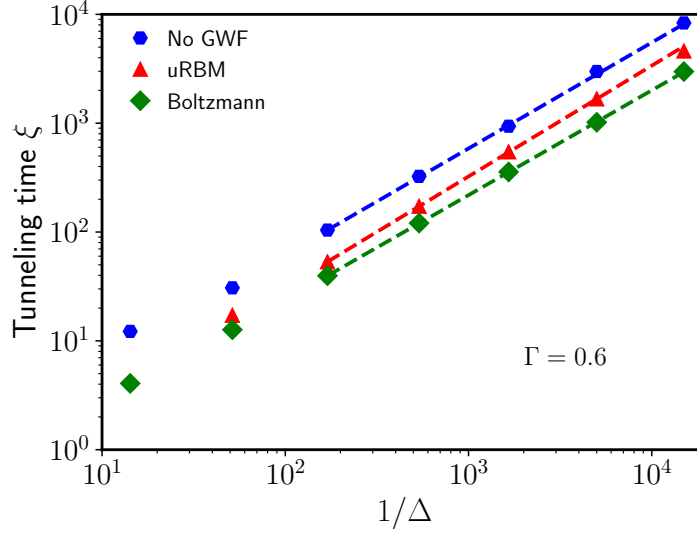


Figure 3.4: Tunneling times ξ of the DMC algorithm implemented without GWF (blue hexagons), with the uRBM GWF (3.47) (red triangles), and with the Boltzmann GWF (3.46) (green diamonds), as a function of the inverse energy gap Δ^{-1} , for the quantum Ising chain at $\Gamma = 0.6$. The dashed lines represent the fitting function $\xi(\Delta) = \alpha\Delta^{-b}$, valid in the large Δ^{-1} regime. In all three cases, the fitted exponent is $b \approx 1$ (see Table 3.2).

quantum Ising chain — possibly down to a polynomial scaling with system size. Unlike the restricted Boltzmann machine originally introduced as a variational ansatz in Ref. [62], the uRBM includes intra-layer interactions. In general, this implies that one cannot analytically trace out the hidden spin configurations.⁸ In order to use uRBMs as GWFs, we employ the extended DMC algorithm described in Ref. [132]. It includes a certain number of additional single-spin Metropolis updates of the hidden spins at every DMC time step. This number has to be made large enough to eliminate spurious correlations among successive walker configuration, which in turn affect the finite- N_w bias. It is quite important to test if and how the possible residual statistical correlations between successive hidden-spin configurations affect the tunneling dynamics.

The tunneling time simulations are performed in the ferromagnetic phase, where the TFIM is characterized by a double-well potential profile. This time we use $p = 10\%$ for the percentage of walkers required to reach negative magnetization. The measurement is repeated about 1 000 times to determine ξ and its standard error. We choose N_w in the range 5 000 – 10 000, which is found to be sufficient to eliminate any systematic error on ξ .

⁸As shown in Ref. [69], the uRBM can be mapped to a constrained matrix product state. In one dimension, this representation allows for an analytical treatment of the hidden degrees of freedom. However, we aim at a general framework that could be applied irrespectively of the dimensionality and the interaction range.

PQMC	α	b
No GWF	0.7(2)	0.97(3)
uRBM	0.32(9)	1.00(2)
Boltzmann	0.28(5)	0.96(3)

Table 3.2: Fitting parameters α and b , describing the small-gap behavior of the DMC tunneling time ξ in the ferromagnetic quantum Ising chain (3.45), according to the fitting function $\xi(\Delta) = \alpha\Delta^{-b}$. The error bars also take into account the fluctuations due to choosing different fitting windows.

The tunneling time ξ obtained with either the Boltzmann or the uRBM GWF as a function of the system size L is qualitatively identical to that without any GWF, see Fig. 3.3. In the large-size regime, where the energy gap Δ is small, the exponential growth of ξ closely matches the scaling of the inverse energy gap $\alpha\Delta^{-1}$, where α is an appropriate prefactor.

The tunneling times obtained with the three methods are plotted in Fig. 3.4 as a function of the inverse energy gap Δ^{-1} for comparison. In all three cases, ξ appears to scale asymptotically linearly with the inverse gap. By fitting the three datasets in the large Δ^{-1} regime with the scaling law $\xi(\Delta) = \alpha\Delta^{-b}$, we obtain the values of the fitting parameters α and b reported in Table 3.2. In all cases, the exponent is consistent with the linear scaling $b = 1$.

3.3.5 The shamrock model

The results for the ferromagnetic Ising chain presented in the previous section indicate that, in an effective double-well system, QMC simulations can efficiently simulate incoherent quantum tunneling and they might therefore be as fast as, or even faster than quantum annealers in solving complex optimization problems. In order to understand if this finding is valid in a more general setup, the authors of Ref. [22] considered a frustrated spin model, described by the Hamiltonian

$$H = -J \sum_{i=0}^{K-1} \sum_{j=2i}^{2i+1} \sigma_{2K}^z \sigma_j^z + (J - \epsilon) \sum_{i=0}^{K-1} \sigma_{2i}^z \sigma_{2i+1}^z - \Gamma \sum_{i=0}^{2K} \sigma_i^x. \quad (3.49)$$

The $L = 2K + 1$ spins are arranged in a clover-shaped bouquet of K rings (see Fig. 3.5), motivating the authors' name *shamrock model* for Eq. (3.49). The first term in the Hamiltonian describes ferromagnetic interactions between the central spin and the outer $L - 1$ spins; the second term couples the two outer spins of each ring through an anti-ferromagnetic interaction which is slightly smaller in modulus than the FM one, $0 < \epsilon \ll J$; and the third term describes a transverse field of intensity Γ , inducing the quantum dynamics.




Figure 3.5: The “shamrock” model of L frustrated spins in a transverse field, Eq. (3.49). It consists of $K = (L - 1)/2$ leaves, each having three spins, one of which is shared among all. The solid dark-green lines represent FM interactions (of strength J) between the central spin and all the others. The dashed light-green lines indicate AFM interactions (of strength $J - \epsilon$) between the outer spins of each leaf. The overall effect is to create 2^K tunneling paths between the degenerate classical ground states.

We investigate the DMC tunneling time in this model using the same protocol as in Section 3.3.4, both without importance sampling and with the Boltzmann GWF. Notice that the classical energy function in Eq. (3.46) now includes the first two terms of Eq. (3.49).

In Fig. 3.6 the tunneling times obtained with the two methods are compared with each other, with the scaling of incoherent quantum tunneling, and with the scaling found in the finite-temperature PIMC simulations of Ref. [22]. In the large $1/\Delta$ regime, the DMC data with GWF are well described by the usual fitting function $\xi(\Delta) = \alpha\Delta^{-b}$, where the fitting parameters are found to be $\alpha = 0.32(7)$ and $b = 1.04(3)$ for the Boltzmann-GWF case. A similar fit, with $b = 0.98(2)$, applies to the no-GWF data. These results indicate that even in a frustrated model such as the shamrock, the DMC tunneling times asymptotically scale with the inverse gap (independently of the chosen GWF), confirming the quadratic speedup compared to incoherent quantum tunneling.

In contrast, we can see that the PIMC results display a faster growth of ξ with the system size, very accurately described by the scaling law $\xi^{\text{PIMC}} \propto 2^K/\Delta^2$, far outpacing both DMC and incoherent quantum tunneling. This pathological slowdown of PIMC simulations was anticipated by the perturbation theory of Ref. [22], which predicts that in frustrated models where two competing ground states are connected by a number of homotopy-inequivalent paths which grows with system size, incoherent quantum tunneling can display a quantum speedup if many inter-path transitions are inhibited by topological obstructions (related to the obstructions discussed in Ref. [119]). The shamrock model was indeed introduced as an example of this scenario, with the PIMC simulations confirming the theoretical prediction even beyond the perturbative regime.




Figure 3.6: Tunneling time ξ in the shamrock model as a functions of the system size L . The DMC results obtained with the Boltzmann GWF (green diamonds) and without GWF (blue pentagons) are compared with the scaling of the incoherent quantum tunneling time $1/\Delta^2$ (black triangles), and with the scaling of the finite-temperature PIMC tunneling time $\xi \sim 2^K/\Delta^2$ [22], where K is the number of leaves in the shamrock. The values of the gap Δ are obtained from exact diagonalization. The model parameters are $\Gamma = 0.5$, $J = 6$, and $\epsilon = 0.2$.

3.3.6 Finite-population bias

As explained in Section 3.3.1, the ground-state energy obtained via DMC simulations is subject to two sources of possible systematic bias, originating from the finite time step $\delta\tau$ (Trotter error) and from the finite random-walker number N_w . The convergence to the $\delta\tau \rightarrow 0$ limit is well-controlled,⁹ and all results presented in this paper have been performed using sufficiently small $\delta\tau$ to make its systematic effect negligible compared to the statistical uncertainty. In this section, we focus on the bias resulting from the finite value of N_w . We limit our discussion to the worst-case scenario, *i.e.*, without importance sampling. We will comment about the role of importance sampling with regard to the systematic error in the conclusions.

We consider the ferromagnetic quantum Ising chain defined in Eq. (3.45). Its ground-state energy per site can be exactly determined via a Jordan–Wigner transformation (Appendix A), obtaining [205]

$$\frac{E_0}{L} = -\frac{J}{\pi} \left(1 + \frac{\Gamma}{J}\right) E(\theta), \quad (3.51)$$

where $E(x)$ is the complete elliptic integral of the second kind and

$$\theta^2 = \frac{4\Gamma J}{(J + \Gamma)^2}. \quad (3.52)$$

This formula is valid in the thermodynamic limit. The results presented in this section have been obtained using sufficiently large system sizes so that finite-size effects are not relevant.

In the main panel of Fig. 3.7a we plot the relative error $e_{\text{rel}} = |E - E_0|/|E_0|$ of the DMC result E with respect to the theoretical value as a function of the transverse field intensity, for different system sizes. These data correspond to a fixed random walker population $N_w = 20\,000$. In the paramagnetic phase $\Gamma > J$, as well as in the $\Gamma \rightarrow 0$ limit, the systematic bias due to finite N_w is negligible. However, in the ferromagnetic phase $0 < \Gamma < J$ a systematic bias is observable, and this bias increases with the system size L . Notice how the maximum point of e_{rel} slightly shifts toward lower values of Γ as L is increased. This is a typical finite-size effect where the transition point of a critical model is slightly offset at finite values of L compared to its true position. In the inset of Fig. 3.7a the relative error is shown as a function of the (inverse) number of walkers $1/N_w$ for a chain of $L = 60$ spins. As expected, in the infinite N_w limit the hardest point in the phase diagram to simulate is

⁹The Trotter error using the simplest factorization, Eq. (3.14), is known to scale as $O(\delta\tau)$, as results from compounding $1/\delta\tau$ steps with $O(\delta\tau^2)$ error each. Faster convergence can be attained by means of more sophisticated factorizations, *e.g.* the symmetric one

$$e^{-\delta\tau(A+B)} = e^{-\delta\tau A/2} e^{-\delta\tau B} e^{-\delta\tau A/2} + O(\delta\tau^3), \quad (3.50)$$

resulting in an $O(\delta\tau^2)$ overall error. It is possibly to go even further [134], but it is seldom useful in practice.

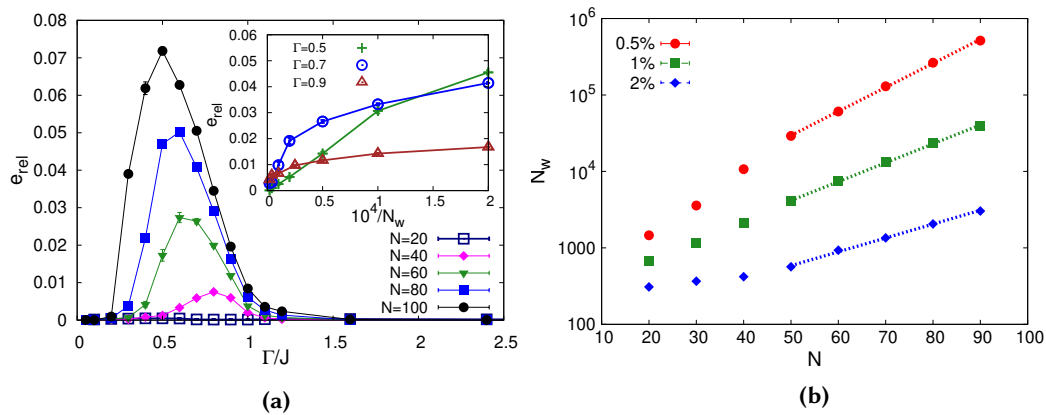


Figure 3.7: (a) Relative error $e_{\text{rel}} = |E - E_0|/|E_0|$ of the DMC result E with respect to the infinite-size theoretical value E_0 (Eq. (3.51)) as a function of the transverse field intensity Γ , for different system sizes L . The average number of random walkers is $N_w = 20\,000$. *Inset:* e_{rel} as a function of the inverse number of walkers $1/N_w$ for different transverse field intensities Γ . The size of the spin chain is $L = 60$. (b) Walker population needed to achieve 0.5%, 1%, and 2% relative error as a function of the system size N . The transverse field intensity is $\Gamma = 0.95J$.

to the left of the critical point, where quantum fluctuations are largest and demand for a very large walker population in order to capture the ground state properties of the system with a reasonable precision.

From Fig. 3.7a we can also see that, unsurprisingly, the number of walkers required to achieve a given relative error increases as the system size gets larger. This dependence is plotted in Fig. 3.7b. In the large- L limit, the data are well described by an exponential fitting function, suggesting that the computational complexity of DMC without importance sampling is exponential in the system size.

3.3.7 Conclusions

We implemented a projective QMC method for quantum Ising models based on the DMC algorithm — in which the transition matrix is defined using a Trotter approximation of the Green’s function — and we investigated the characteristic time of tunneling events in problems characterized by an effective double-well energy landscape. We found that the DMC tunneling time increases with the system size as the inverse of the gap, that is, more favorably than the incoherent tunneling time, in all models we considered. These included a continuous-space double-well model, the ferromagnetic quantum Ising chain and the more challenging shamrock model with frustrated couplings. This scaling is to be contrasted with previous studies based on finite-temperature PIMC simulations, where topological obstructions in the shamrock model were found to cause a considerable

(exponential in the system size) slowdown of the PIMC tunneling time.

We have also investigated how guiding wave functions affect the tunneling dynamics of DMC simulations, using two approximate variational ansatzes as well as the numerically-exact ground-state wave function (for the 1D model). Remarkably, for all GWFs we find a linear relation between tunneling rate and ground energy gap in the asymptotic regime of large tunneling times, corresponding to a high potential barrier in the double well, or to large system sizes in the spin models. The semiclassical theory we provided explains this linear relation in the case of double-well-type potentials when the exact ground-state wave function is chosen as GWF. It is worth stressing that this linear relation represents a quadratic speedup compared to the expected tunneling rate of a physical quantum annealer. The proof we presented relies on the local validity of the semiclassical approximation for the ground-state wave function. It is an interesting challenge to try to formulate a more general derivation which does not invoke WKB theory, or to instead exhibit a counterexample where a violation of the $1/\Delta$ scaling may be observed.

Analyzing if and to what extent QMC algorithms can efficiently simulate quantum tunneling is of critical importance to understand if quantum annealing devices can outperform classical optimization methods. At first glance, the exponential memory requirement found in Section 3.3.6 for fixed-precision measurements is enough to make large-size DMC simulations unfeasible. The situation changes, however, when a physically motivated ansatz is available for the importance sampling of the ground-state wave function. Indeed, it is known that the use of GWFs can drastically improve the complexity scaling of the algorithm, to the point that it was suggested in the case of an Ising chain that the computational cost may possibly change from exponential to polynomial when good GWFs are used [132]. Our findings indicate that projective QMC simulations performed with accurate GWFs may allow one to efficiently simulate both the equilibrium ground-state properties as well as the tunneling dynamics of quantum annealers. Therefore, they may be used as a relevant benchmark in the development of novel quantum annealing devices, and they represent a promising quantum-inspired optimization algorithm.

Clearly, more challenging models should be addressed to further benchmark the efficiency of the DMC tunneling dynamics. Relevant test beds could be Ising spin glasses in higher dimensions, where one expects accurate variational ansatzes to be less easily available. Suitable candidates are restricted [62, 178] and unrestricted Boltzmann machines [38, 132], both of which have been shown to be amenable to be used as GWFs in projective QMC simulations [207]. Deeper neural network ansatzes, *e.g.* the deep convolutional neural networks of Ref. [67] or the recurrent neural networks of Ref. [63], might also be adopted.

Chapter 4

Population Transfer

Πολλοὶ γάρ εἰσιν κλητοί, ὀλίγοι δὲ ἐκλεκτοί.

For many are called, but few are chosen.

Matthew 22:14 (DRA)

4.1 Energy matching and the population transfer algorithm

In this chapter we return to the class of combinatorial optimization problems and the possibility of exploiting quantum effects in order to devise more efficient strategies for tackling them. However, we introduce a twist: rather than trying to find low-cost assignments *ex novo*, we are now interested in the computational problem of finding assignments with similar cost to another one, which is given beforehand.

Recall that a combinatorial optimization problem (COP) is defined by a real-valued *cost function* (or Hamiltonian) H_p defined on some n -variable set $\mathcal{D} \subset \mathcal{X}^n$, where \mathcal{X} has finite cardinality. In particular, without loss of generality we may consider $\mathcal{X} = \{0, 1\}$, so that the cost function is defined on (a subset of) the *Boolean hypercube* of all possible n -bit words z . The objective is to minimize H_p over its domain \mathcal{D} .¹ In other words, in optimization we want to find an explicit assignment z_0 such that $H_p(z) \geq H_p(z_0)$ for all $z \in \mathcal{D}$. Such a z_0 is called a *solution* to the optimization problem. Given a COP, there is a family of associated decision problems asking whether, for a given E_0 , there exists any $z_0 \in \mathcal{D}$ such that $H_p(z_0) \leq E_0$. If each of them is in the NP class (*i.e.* if the cost function is polynomially-computable) the optimization problem is said to belong to NPO [1, 29].

¹In many settings, one wants instead to *maximize* some target function F (*e.g.* fitness in the context of a genetic algorithm), but this is of course equivalent to minimizing the cost function $-F$.

The solution of optimization problems is generally a very hard problem, and *approximation algorithms* are typically employed to find not a strict solution of the problem, but rather an element of \mathcal{D} whose cost is guaranteed to be closer to the optimal one than a given multiplicative error (*i.e.* an *approximate solution*). In some cases, this can substantially decrease the complexity of the original problem: in particular, optimization problems in the APX class (a subset of NPO) admit polynomial-time approximation algorithms for some fixed multiplicative error, in the sense that there is a $c > 0$ such that it only takes time $O(\text{poly}(n))$ to produce an assignment whose cost is at most $(1 + c)$ times the optimal cost. For some problems, including the Traveling Salesman Problem in Euclidean space [23], this even holds for *any* fixed c (PTAS class).

However, a result known as the PCP theorem [24] implies that it is impossible (unless $P = NP$) to efficiently find approximate solutions to MAX-3-SAT with a better multiplicative error than the one valid for *random* assignments. In other words, problems such as MAX- k -SAT ($k \geq 3$) are hard to solve even approximately.

In this chapter, we consider the situation where one already has in his hands an (approximate) solution of a hard optimization problem, and is interested in the problem of finding *another* one. More precisely, we define the following computational primitive:

Definition 4.1.1. (Energy matching) The $(g(n), H_p, z_0)$ -*energy matching problem* consists of the following:

given: a cost function $H_p : \{0, 1\}^n \rightarrow \mathbb{R}$ over n binary variables and an n -bit string $z_0 \in \{0, 1\}^n$ with $H_p(z_0) = E$;

find: an n -bit string $z \neq z_0$ such that $E(1 - g(n)) \leq H_p(z) \leq E(1 + g(n))$.

In Def. 4.1.1, E is the *target energy* and $g(n)$ defines the approximation one is willing to accept in the solution. For genuine matching, one wants $g(n) \xrightarrow{n \rightarrow \infty} 0$. If one defines an “energy shell” around E of width $2\Delta E$,

$$\Omega(E, 2\Delta E) = \{z \mid E - \Delta E \leq H_p(z) \leq E + \Delta E\}$$

(cf. Eq. (1.8)), then the *target states* for the energy matching problem are the states in the “punctured” energy shell, $z \in \Omega_n = \Omega(E, 2g(n)) \setminus \{z_0\}$.

The name “energy matching” is clearly motivated by physics, where H_p is a “problem Hamiltonian” in the same spirit of the QA algorithm (cf. Section 1.3.2). However, its scope and usefulness are not confined to physics — indeed, they are completely general. Consider any optimization problem, and suppose you have an algorithm that is able to provide you with a solution² to said problem. In a realistic setting, it is often the case that a single solution is not enough. One may find that the purported solution has some

²Here and for the rest of the paragraph, we use “solution” as a shorthand for “approximate solution”.

undesirable property, *e.g.* it may violate some constraints that were not dutifully embedded in the cost function, making it unviable. If the approximation scheme that generated the solution is extremely costly, or if it is biased in a way that will have it always output the same solution or slight variations thereof (*e.g.* if it is completely deterministic, so that it always produces the same output given identical inputs), the need may arise to find new, fundamentally different solutions. This is nothing but the problem of energy matching.

The requirement to produce *sufficiently distinct* solutions from the initial one can be defined in a rigorous way by augmenting the energy requirement with a lower bound on the *Hamming distance* (minimum number of bit flips) between the initial and final solutions:

Definition 4.1.2. (Energy matching with distance bound) The $(g(n), d(n), H_p, z_0)$ -energy matching problem is the $(g(n), H_p, z_0)$ -energy matching problem (Definition 4.1.1) subject to the additional constraint

$$|z - z_0|_H > d(n) \quad (4.1)$$

on the set of acceptable solutions,³ where $|z|_H$ is the Hamming weight (number of nonzero bits) of n -bit string z and $z - z_0$ represents bit-wise subtraction modulo 2 (namely, the XOR ^{$\otimes n$} operator).

The lower bound on the distance, introduced to exclude “trivial solutions” which are only slight perturbations of the initial state, makes it harder for a local-search heuristic to solve the problem, as a volume which is exponentially large in $d(n)$ must now be explored. For example, it is known [31] that the decision version of the energy matching problem for a ground state of MAX- k -SAT is NP-complete when $d(n) = n/3$.

The *population transfer* (PT) algorithm is a proposed strategy to solve the energy matching problem exploiting the properties of quantum mechanics. The term “population transfer”, used since the 1980s in the quantum optics community to refer to the pumping of photons from one energy level to another (see *e.g.* [64]), was first applied to the specific scenario described hereafter by the authors of Ref. [234], although the algorithm itself was proposed some time earlier by Baldwin and Laumann [31].

It is conceptually very simple: suppose that the cost function can be encoded into a spin Hamiltonian H_p , diagonal in the σ^z basis, and that an external transverse driver is available to build the full Hamiltonian

$$H = H_p(\{\sigma_i^z\}) - \Gamma \sum_{i=1}^n \sigma_i^x. \quad (4.2)$$

This is reminiscent of the standard QA setup described in Section 1.3.2 (cf. Eq. (1.52)), except that there is no annealing schedule involved. Rather, a quantum quench protocol is considered, as in Eq. (2.14). More precisely, we have the following:

³Notice that Definition 4.1.1 corresponds to the particular case $d(n) = 0$.




Figure 4.1: Schematic representation of the population transfer (PT) algorithm described in Definition 4.1.3. The circuit is iterated until a state $|z\rangle$ turns up satisfying the $(g(n), d(n))$ -energy matching conditions.

Definition 4.1.3. (Population transfer) The (Γ, t_{fin}) -population transfer algorithm applied to the $(g(n), d(n), H_p, z_0)$ -energy matching problem proceeds as follows:

- 1) prepare⁴ the state $|z_0\rangle$ and the Hamiltonian (4.2) with the prescribed value of Γ ; call $E_{z_0} = \langle z_0 | H_p | z_0 \rangle$;
- 2) evolve the system according to the propagator generated by H , obtaining the state $|\psi(t_{\text{fin}})\rangle = e^{-iHt_{\text{fin}}} |z_0\rangle$;
- 3) perform a projective measurement (readoff) in the computational basis and thereby obtain a bit string z with probability

$$p(z) = |\langle z | \psi(t_{\text{fin}}) \rangle|^2;$$

- 4) iterate steps 2 and 3 until the energy $E_z = \langle z | H_p | z \rangle$ satisfies the matching condition

$$|E_z - E_{z_0}| \leq g(n) |E_{z_0}|$$

while simultaneously

$$|z - z_0|_{\text{H}} > d(n).$$

See Fig. 4.1 for a graphical representation of steps 1–3 of the algorithm.

As one may glean from Def. 4.1.3, the efficacy of PT on solving the energy matching problem critically depends on the choice of the two user-defined parameters Γ and t_{fin} characterizing the transfer dynamics. In particular, the transverse field should be sufficiently strong to promote PT, but not so strong as to excite the initial state way out of the target shell Ω_n ; likewise, t_{fin} should be taken large enough for transfer to occur, but minimally so in order not to waste computation time. The problem of finding good values for Γ and t_{fin} will become clearer in the course of the chapter. In particular, one of the main goals of Section 4.3 will be to study how this choice of parameters can be made properly for a specific model. Before attacking this problem, however, we present a simpler case where analytical calculations can be carried quite far, yielding precise answers in the large- n regime.

⁴Notice that the state $|z_0\rangle$ is, at least in principle, exactly known (as it is a computational-basis state) and easy to prepare. In particular, there is no cloning involved in running steps 2 and 3 of the algorithm multiple times until the desired solution is found.

4.2 PT in the impurity band model

We start our discussion of PT with the model studied by Smelyanskiy *et al.* [234], which the authors call the *impurity band model*. The problem Hamiltonian H_p is taken to act as the zero operator on most of the Hilbert basis, except for a “small” (but still exponential-in- n) number of states, which are instead assigned an extensive energy $E_z = -n + \epsilon_z$, where the uniformly random fluctuations $\epsilon_z \in [-W/2, W/2]$ are independent and bounded within a sub-extensive interval, $W = o(n)$.

4.2.1 Model definition and motivation

We now proceed to give a more precise definition of the model.

Notation. We use $a, b, \dots \in [D] - 1 = \{0, \dots, D - 1\}$ to label the computational basis, where $D = 2^n$. Then $z_a = a_{n-1} \dots a_1 a_0$ denotes the n -bit string representing integer a in binary. z, z', \dots are used for generic n -bit strings when integer labeling does not matter. We switch to upper case (A, B, \dots) for the labeling of impurity band states. These indices vary within $\mathcal{M} \subset [D] - 1$. We reserve i, j, \dots for spatial indexing, running over $[n] - 1$.

We omit the range of variability of indices whenever they are implied to run over their entire domain. Therefore \sum_a means $\sum_{a=0}^{D-1}$, \sum_A means $\sum_{A \in \mathcal{M}}$, *etc.*

Definition 4.2.1. (Impurity band model) The (μ, W) -*impurity band model* (IBM), $0 \leq \mu \leq 1$, is defined as the ensemble of Hamiltonians, diagonal on the z -basis $\{|z_a\rangle\}_a$, with action

$$H_p |z\rangle = \mathcal{E}(z) |z\rangle, \quad (4.3)$$

where

$$\mathcal{E}(z_a) = \begin{cases} -n + \epsilon_{z_a} & \text{if } a \in \mathcal{M}, \\ 0 & \text{else,} \end{cases} \quad (4.4)$$

\mathcal{M} is a uniformly random subset of $[D] - 1$ with cardinality

$$|\mathcal{M}| = M = \lfloor 2^{\mu n} \rfloor \quad (4.5)$$

and ϵ_{z_A} is a quenched random variable uniformly picked from the interval $[-W/2, W/2]$ for each $A \in \mathcal{M}$.

This peculiar choice for $\mathcal{E}(z)$ is meant to replicate the features of typical “golf-course” energy landscapes, with most states lying on a constant-energy-density *plateau* and a few low- and high-lying states interspersed in a quasi-random fashion among the rest.

Obviously, in order for this to hold we need $W \ll n$. Moreover, for technical reasons to be mentioned later, we will always consider the regime $\mu \ll 1$. Therefore, the *impurity band* (IB) \mathcal{M} is going to be both narrow and sparse.

Definition 4.2.2. (Quantum impurity band model) The (μ, W, Γ) -quantum impurity band model (QIBM) is obtained by equipping the (μ, W) -IBM with the transverse driver

$$H_d = -\Gamma \sum_i \sigma_i^x, \quad (4.6)$$

where σ_i^x acts as the Pauli-X operator on qubit i and the identity elsewhere. The driver can also be written in term of the global spin's x -component, $H_d = -2\Gamma S^x$.

As we are only interested in the low-energy dynamics, *i.e.* the dynamics restricted to the impurity band \mathcal{M} , we are going to switch to a simpler, effective description of the system. To this end, we make use of the following, easy-to-check property of the resolvent $G(E) = (E - H_d)^{-1}$:

$$G(E)H_p |\psi\rangle = |\psi\rangle \quad (4.7)$$

where $H |\psi\rangle = E |\psi\rangle$. Pre-multiplying this by $\sqrt{H_p}$ and defining the (non-normalized!) projected state $|\Psi\rangle = \sqrt{H_p} |\psi\rangle$, simple algebra leads us to the matrix identity

$$(H_p + \Lambda) |\Psi\rangle = E |\Psi\rangle, \quad (4.8)$$

where $\Lambda = \sqrt{H_p} H_d G(E) \sqrt{H_p}$.

Now, projecting onto the IB through $\langle z_A |$ and using the fact that H_p (and therefore $\sqrt{H_p}$) annihilates all states outside the IB, so that we can insert a resolution of the identity between $G(E)$ and $\sqrt{H_p}$ and simplify it from $\sum_a |z_a\rangle \langle z_a|$ to $\sum_A |z_A\rangle \langle z_A|$, we get the following “non-linear eigenvalue equation” for $|\Psi\rangle$:

$$\sum_B \mathcal{H}_{AB}(E) \Psi_B = E \Psi_A, \quad (4.9)$$

where

$$\begin{cases} \mathcal{H}_{AB}(E) = \delta_{AB} \mathcal{E}(z_A) + \sqrt{\mathcal{E}(z_A) \mathcal{E}(z_B)} c_{AB}(E), \\ c_{AB}(E) = -\langle z_A | H_d G(E) | z_B \rangle. \end{cases} \quad (4.10)$$

Notice that the $c_{AB}(E)$ coefficient is a property of the driver alone, and does not depend on H_p . However, deriving Eq. (4.9) from (4.8) required us to use the “projector” property of H_p , which is very nongeneric. Indeed, the IBM was chosen by the authors specifically so that this simplification would be available. For a model such as the QREM, where there is no clear-cut distinction between “band states” and “plateau states”, Eq. (4.9) could only be interpreted as an approximation.

4.2.2 Semiclassical treatment and perturbation theory

The computation of $c_{AB}(E)$ can be carried out in the semiclassical (WKB) approximation. We only sketch the starting point, and refer to Ref. [234] for the full calculation. The WKB result improves on simpler perturbative estimations obtained via the forward-scattering approximation [159] by an exponential prefactor.

First, notice that the $c_{AB}(E)$ coefficient only depends on the Hamming distance $r = |z_A - z_B|_H$. This can be seen by explicitly writing the coefficient in the x -basis,

$$\begin{aligned} c_{AB}(E) &= \sum_{x,x'} \langle z_A|x \rangle \langle x|H_d(E - H_d)^{-1}|x' \rangle \langle x'|z_B \rangle \\ &= \frac{-\Gamma}{2^n} \sum_x \frac{(-1)^{x \cdot (z_A - z_B)} (n - 2|x|_H)}{E + \Gamma(n - 2|x|_H)} \\ &= \frac{1}{2^n} \sum_{m=0}^{n-r} \sum_{l=0}^r \binom{n-r}{m} \binom{r}{l} \frac{(-1)^l}{1 + \frac{E}{\Gamma(n-2(l+m))}}, \end{aligned} \quad (4.11)$$

where in the second passage we used the overlap $\langle x|z \rangle = (-1)^{x \cdot z} 2^{-n/2}$ (with the dot product notation $x \cdot z = \sum_i x_i z_i$) and the diagonal action of the driver on the x -basis, $\langle x|f(H_d)|x' \rangle = f(-\Gamma(n - 2|x|_H)) \delta_{xx'}$; in the last passage, we split x into its $n - r$ bits at the positions where z_A and z_B agree plus the r bits where they disagree, and then populated these substrings in all possible ways with m and l '1's, respectively, so that $(-1)^{x \cdot (z_A - z_B)} = (-1)^l$ and $|x|_H = l + m$.

Then, rewrite Eq. (4.10) in the form

$$c(E, r) = \langle z_0 | 1 + \frac{E}{E + 2\Gamma S^x} |z(r) \rangle, \quad (4.12)$$

where without loss of generality we bracket between $|z_0 \rangle = |0 \cdots 0 \rangle$ and any state $|z(r) \rangle$ with exactly r spin-downs, e.g. $|z_{2^r-1} \rangle = |0^{n-r} 1^r \rangle$. Now, since the result must be the same regardless of the positioning of the r spin-downs, it is also equal to the average over such choices. But the (normalized) average state is

$$\left| \left\langle \frac{n}{2} - r; \frac{n}{2} \right\rangle \right\rangle = \frac{1}{\sqrt{\binom{n}{r}}} \sum_{|z|_H=r} |z \rangle, \quad (4.13)$$

namely the simultaneous eigenstate of S^z and S^2 with respective eigenvalues $M = \frac{n}{2} - r$ and $S(S+1) = \frac{n}{2}(\frac{n}{2} + 1)$. This leads us to rewriting Eq. (4.12) as

$$c(E, r) = \delta_{r,0} - \frac{E}{\sqrt{\binom{n}{r}}} G_{n/2-r, n/2}(E), \quad (4.14)$$

with the Green function defined between maximal-spin states:

$$G_{M,n/2}(E) = \left\langle \left(M; \frac{n}{2} \right) \left| \frac{1}{E + 2\Gamma S^x} \right| \left(\frac{n}{2}; \frac{n}{2} \right) \right\rangle. \quad (4.15)$$

As the spin is extensively large, a WKB treatment is justified, starting from the substitution

$$G_{M,n/2}(E) \sim \exp \left[i \int_0^M p(k, E) dk \right], \quad (4.16)$$

with the appropriate semiclassical momentum $p(k, E)$. After tedious calculations, one reaches the following expression for $c(E, r)$, valid in the vicinity of $r = n/2$:⁵

$$c(E, r) \propto \frac{1}{\sqrt{\binom{n}{r}}} e^{-\theta(\Gamma)} \sin \phi(E, r), \quad (4.17)$$

where $\theta(\Gamma) \sim 1/(4\Gamma^2)$ at large Γ .

The next step of the authors of Ref. [234] is to treat both the IB width W and the off-diagonal propagator coefficients $c_{A \neq B}(E, r)$ as perturbations to the zero-order Hamiltonian

$$H_{AB}^{(0)}(E) = \delta_{AB} n(c(E, 0) - 1). \quad (4.18)$$

This is justified by choosing $W \ll n$, as well as noticing that the off-diagonal propagator coefficients $c_{A \neq B}$ are suppressed with respect to c_{AA} by an exponentially small factor $\binom{n}{r}^{-1/2}$. The first perturbative correction is then given by the so-called *downfolded Hamiltonian*

$$\Delta^{(1)} H_{AB} = \delta_{AB} \epsilon_{z_A} + (1 - \delta_{AB}) \sqrt{2} V(r) \sin \phi(r) \quad (4.19)$$

where ϵ_z is the energy of z relative to the IB center (cf. Eq. (4.4)), $r = |z_A - z_B|_H$ and $V(r)$ is a complicated function of r representing the off-diagonal WKB matrix element between states $|z_A\rangle$ and $|z_B\rangle$, up to an oscillating factor $\sin \phi(r)$ also computed via WKB theory.

The final step is to calculate the pdf of the off-diagonal elements, or rather their squares $(\Delta^{(1)} H_{A \neq B})^2$, as induced by the random distribution of distances between IB states. The reason for focusing on the squares is that our ultimate goal is to find the distribution of the time required to achieve population transfer from an initial state $|z\rangle$, and this can be estimated using the Fermi golden rule, which notoriously involves the squared matrix elements of the perturbation $\Delta^{(1)} H$.

When squaring the off-diagonal part of Eq. (4.19), the oscillatory term can be shown to merely provide a factor 1/2 under disorder average, while an application of (a variant

⁵Technically, one must also ensure that the value of Γ is sufficiently far from one of the resonant points $\Gamma_p \approx n/(n - 2p)$ ($p \in \{0, \dots, n\} \setminus \{n/2\}$) so that the IB center does not coincide with one of the unperturbed levels of H_d . In that case, the classical and driver parts are resonant and perturbation theory breaks down. At large enough n , the resonant regions are exponentially thin and Γ can always be taken to be nonresonant.

of) the central limit theorem leads to the determination of the distribution of $V(r)^2$, which turns out to be fat-tailed: in particular, by writing $V^2 = V_{\text{typ}}^2 w$ with V_{typ} the typical⁶ value of the matrix element, w is distributed according to⁷

$$p(w) = \frac{1}{w^2 \sqrt{\pi \log w}}. \quad (4.20)$$

We can then see that the ensemble Eq. (4.19) consists of $M \times M$ matrices with $O(V_{\text{typ}})$ off-diagonal Lévy-distributed entries and an $O(W)$ uniformly random diagonal. Following the authors, we refer to it as the preferred-basis Lévy matrix (PBLM) ensemble. It is reminiscent of the so-called Rosenzweig–Porter (RP) ensemble [220], where all entries are Gaussian-distributed, like in the GOE, but off-diagonal elements have parametrically smaller variances. The RP ensemble is known to display three different regimes depending on the ratio of these variances, which suggests something similar might be expected of the PT dynamics of the IBM. It is worthwhile to state the facts with a higher degree of precision.

Definition 4.2.3. (Rosenzweig–Porter model) The γ -RPM is the ensemble of real symmetric $D \times D$ Hamiltonians whose diagonal entries H_{aa} are i.i.d. random variables taken from the standard normal distribution $\mathcal{N}(0, 1)$, and whose upper-triangle entries $H_{a>b}$ are i.i.d. random variables taken from $\mathcal{N}(0, D^{-\gamma/2})$ (i.e. normal variables with zero mean and variance $D^{-\gamma}$).

This model has received some recent attention due to its close relationship with the Anderson model on regular random graphs (RRG), which in turn is linked to the idea of MBL as a form of Anderson localization in the Fock space. The following result has been established numerically by Kravtsov *et al.* [153] and later given more rigorous foundations by Facchetti, Vivo and Biroli [96].⁸

Theorem 4.2.1. (Spectral regimes in the RPM) All eigenstates of the γ -RPM are:

- localized on a single site, in the sense that

$$|\psi_a^{(k)}|^2 = \delta_{a, a_0(k)} + O(D^{-1}) \quad \text{for some } a_0(k), \quad (4.21)$$

if $\gamma \geq 2$;

⁶As the typical Hamming distance between states is $r \sim n/2$, for large n the typical value of $V(r)$ coincides with the most likely value, $V_{\text{typ}} \sim V(n/2)$.

⁷The following form for $p(w)$ is only valid for $\mu \ll 1$, motivating the choice of taking the sparse-band regime.

⁸In both cases, the authors actually work with a generalized version of the RPM where diagonal elements may have a non-Gaussian distribution.

- *non-ergodic extended (NEE), namely spread over $O(D^\alpha)$ sites with $\alpha = 2 - \gamma \in (0, 1)$, i.e.*

$$|\psi_a^{(k)}|^2 \sim \begin{cases} D^{-\alpha} & \text{if } a \in \mathcal{V}_k \subset [D] - 1 \text{ with } |\mathcal{V}_k| \sim D^\alpha \\ o(D^{-1}) & \text{otherwise,} \end{cases} \quad (4.22)$$

if $1 < \gamma < 2$;

- *completely delocalized, i.e.*

$$|\psi_a^{(k)}|^2 \sim D^{-1} \quad \text{for all } a, \quad (4.23)$$

if $0 \leq \gamma \leq 1$.

Moreover, it is known that the structure of the RP spectrum in the NEE regime consists of multiple distinct *minibands* of states. Each miniband is polynomially (in D) narrower than the full spectral width and contains states whose supports overlap considerably, while being essentially nonoverlapping with those of the states from other minibands. We can think of the minibands as “communities” of eigenstates, each living for the most part on the same territory which is disjoint from that of all other communities.⁹ The ergodic transition happens when all the communities merge together into a single country, whereas the onset of the fully-localized regime corresponds to each community shrinking down to a single household.

One can show numerically that a similar phenomenon occurs in the PBLM ensemble depending on the control parameter $W/(MV_{\text{typ}}) = \lambda M^{\gamma/2-1}$, where λ is an $O(M^0)$ constant and we chose the following parametrization of W :

$$\frac{W}{V_{\text{typ}}} = M^{\gamma/2}. \quad (4.24)$$

With this definition of γ , the three regimes are actually determined by the same conditions of Theorem 4.2.1. As we will discuss in the next section, this has a profound consequence on the PT dynamics of the model.

4.2.3 PT time and quantum speedup

Smelyanskiy *et al.* follow up with an approximate treatment of the downfolded dynamics by treating the PT problem as a Fano–Anderson model, where an initial discrete state is allowed to decay into a band. This can only be done in the extended phase $\gamma < 2$, where a “continuum” band exists for the initial state to decay into. In the $\gamma \geq 2$ regime, the initial state is isolated from the rest of the Hilbert space and no decay occurs.

⁹Note however that in the IBM these territories are really collections of isolated points; in more realistic models with spatially correlated potentials, the territories are generally comprised of collections of connected clusters in the Boolean cube.

The Fano–Anderson model is the simplest Hamiltonian one can write for this kind of systems: a diagonal part containing the unperturbed energies of all available states, plus a minimal coupling between the initial state and the decay states. The energies of the decay states are augmented by a small imaginary part to smoothen out finite-size effects, and the coupling coefficients between initial and decay states are given by the downfolded matrix elements Eq. (4.19), *i.e.* they are perturbatively determined.

This model can be solved for the initial state's wave function, whose rate of decay is known to be determined (to leading order in $\Delta^{(1)}H$) by the imaginary part of the self-energy:

$$\begin{aligned}\psi_A(z_A, t) &= \int_{-\infty}^{\infty} \frac{dz}{\pi} \frac{\Sigma_A''(z)e^{-izt}}{(z - \Sigma_A'(z) - \epsilon_A)^2 + (\Sigma_A''(z))^2} \\ &\sim \exp \left[-i(\epsilon_A + \Sigma_A'(\epsilon_A))t - \Sigma_A''(\epsilon_A)t \right],\end{aligned}\quad (4.25)$$

where the self-energy is defined by

$$\Sigma_A(z) = \Sigma_A'(z) - i\Sigma_A''(z) = \sum_{B \in \mathcal{M} \setminus \{A\}} \frac{(\Delta^{(1)}H_{AB})^2}{z - \epsilon_B + i\eta} \quad (4.26)$$

with η much larger than the level spacing but much smaller than the decay rate $\Sigma_A''(\epsilon_A)$ (a condition to be verified self-consistently).

Using the Plemelj–Sokhotski theorem and plugging in Eq. (4.19),

$$\begin{aligned}\Sigma_A''(\epsilon_A) &= \pi \sum_{B \in \mathcal{M} \setminus \{A\}} \left(\Delta^{(1)}H_{AB} \right)^2 \delta_\eta(\epsilon_A - \epsilon_B) \\ &= 2\pi \sum_{r=1}^n V^2(r) \sin^2 \phi(r) \left(\sum_{B \in \mathcal{M} \setminus \{A\}} \delta_\eta(\epsilon_A - \epsilon_B) \delta_{r, |z_A - z_B|_{\mathbb{H}}} \right),\end{aligned}\quad (4.27)$$

with δ_η the zero-mean Cauchy distribution of width η . In the last step, the self-energy was split into its n different channels corresponding to all possible distances from A to the other available states.

The parenthesized expression, which is just the density $\rho_\eta^A(r)$ of IB states at a given distance r from $|z_A\rangle$, can be approximated as $\rho_\eta^A(r) \approx p(\epsilon_A)M_A(r)$, with

$$M_A(r) = \sum_B \delta_{r, |z_A - z_B|_{\mathbb{H}}} \sim \binom{n}{r} \quad (4.28)$$

the number of IB states at distance r from $|z_A\rangle$, whose fluctuations we can neglect when taking the disorder average. Finally, combining with Eq. (4.20) and the fact that weakly-correlated Lévy random variables sum to a stable-distributed variable, one finds that the

decay rate $\tau^{-1} = 2\Sigma''_a(\epsilon_a)$ can be written as

$$\frac{1}{\tau} \stackrel{d}{=} \frac{\pi V_{\text{typ}}^2}{W/M} (\sigma_M x + b_M) \quad (4.29)$$

with

$$b_M \sim \sqrt{\log M}, \quad \sigma_M \sim \frac{1}{\sqrt{\log M}} \quad (4.30)$$

and x distributed according to

$$p(x) = \int_{-\infty}^{+\infty} \frac{dt}{2\pi} \exp \left[-itx - \frac{2i}{\pi} t \log |t| - |t| \right], \quad (4.31)$$

which is the so-called *stable distribution* with stability, skewness, scale and location parameters $(\alpha, \beta, s, \ell) = (1, 1, 1, 0)$, respectively. It is also named the Landau distribution, as it was first used by him to describe the energy loss of particles across a thick medium [154]. This distribution asymptotically behaves as $p(x) \sim 2/\pi x^2$ ($x \gg 1$).

The reason why the decay rate has such heavy tails is that its value is essentially determined by the distance between the initial state and its closest resonance, for which the matrix element $(\Delta^{(1)} H_{ab})^2$ is maximal. When the distance is anomalously smaller than the typical value $d = n/2$, the decay rate will be correspondingly higher — a typical large deviation effect giving rise to Lévy-type distributions.

The authors then go on to improve this approximation scheme using a cavity method in the spirit of the Abou-Chakra–Anderson–Thouless equations [5] (though with added complications from nonlinearities in the self-energy), again finding a stable distribution for τ^{-1} , but with the difference that the $\log M$ in Eqs. (4.30) is to be replaced with $\log \Omega$, where Ω is the number of states in the *miniband* where z_A lives (as opposed to the full impurity band). This number can be estimated to be

$$\Omega \sim M^{2-\gamma}, \quad (4.32)$$

so in particular $\Omega \ll M$ in the NEE phase; the logarithmic dependence means that σ_Ω is merely rescaled by a prefactor $\sqrt{2-\gamma}$.

We are finally ready to give an estimate of the *typical* PT time. Notice that the Landau distribution has infinite mean, so a meaningful way to define typicality must use a different statistics; the easiest way is to use the mode (the scaling would not change by using the median instead). The mode is obtained by simply setting $x = 0$ in Eq. (4.29), which yields

$$\tau_{\text{typ}} = \frac{1}{\pi} \left(\frac{W}{V_{\text{typ}}} \right) \frac{1}{V_{\text{typ}} M b_M} \sim \frac{1}{V_{\text{typ}} \sqrt{\Omega \log \Omega}}, \quad (4.33)$$

where Eqs. (4.24), (4.30) and (4.32) were used in the last passage. The value of V_{typ} can be obtained by the WKB calculation for the downfolded Hamiltonian, revealing an

asymptotic scaling $V_{\text{typ}} \sim 2^{-n/2} e^{-n\theta(\Gamma)}$ up to subexponential corrections. This means that, to exponential leading order in n , we have

$$t_{\text{PT}} = \tau_{\text{typ}} \sim \sqrt{\frac{2^n}{\Omega}} e^{2\theta(\Gamma)n}. \quad (4.34)$$

The expression under square root corresponds to the classical time $t_{\text{cl}} = 2^n/\Omega$ required to search a database of 2^n entries for any of Ω marked items. Its square root is then equivalent to the time complexity of a Grover search [116] on the same database. We can see that the PT time is augmented by an additional factor $e^{2\theta(\Gamma)n}$, which means that the quantum speedup is not quite quadratic but rather has a “gain exponent” of

$$\nu = \frac{\log t_{\text{cl}}}{\log t_{\text{PT}}} \sim \frac{2}{1 - \frac{4\theta(\Gamma)}{[1-(2-\gamma)\mu]\log 2}} \sim 2 \left(1 + \frac{1}{C_\gamma \Gamma^2} \right) \quad (4.35)$$

compared to the classical time, in the sense that $t_{\text{PT}} \sim t_{\text{cl}}^{1/\nu}$. The rightmost expression in Eq. (4.35) is valid in the large- Γ limit, with the constant $C_\gamma \approx \log 2$. Recall that μ is the (small) parameter determining the size of the IB, Eq. (4.5), while γ parametrizes its width, Eq. (4.24). Remarkably, the gain exponent can be made arbitrarily close to the Grover value 2 by simply increasing the transverse field.

The main advantage of using PT over Grover is that the latter algorithm requires fine-tuning of the external field with exponential accuracy, when initializing the system in the standard fully-symmetrized state of the computational basis (*i.e.* the ground state of the transverse driver). The reason is that the success of the analog Grover algorithm [101,217] hinges on projecting the time-evolved state on the desired solution manifold after a Rabi oscillation from the initial state. This manifold is exponentially small compared to the entire Hilbert space, and one must be comparatively precise in the determination of the measurement time (or equivalently, of the transverse field) in order to achieve the desired resonance. On the contrary, population transfer relies on a multi-channel process where numerous virtual paths with $1 \leq r \leq n$ contribute non-negligibly to the final probability (this is the case because the growth of $M_A(r)$ with r exactly compensates the simultaneous decrease of $V(r)$). This allows PT to be very robust against errors in Γ , which does not need to be fine-tuned but merely “large”.

4.3 PT in the random energy model

In this section, we tackle a slightly different toy model with a more realistic density of states, the *quantum random energy model*. It can be thought of as a Gaussian version of the QIBM described in the previous section, with the difference that there is no clear-cut divide between “impurity” and “plateau” states. Instead, a continuous system of levels

connects the low- and high-temperature regions of the spectrum. Even this seemingly minor modification makes it nontrivial to apply the formalism of the previous section to this case, although the model remains somewhat analytically tractable, to the point that several [31,32,98,235] theoretical treatments have been proposed, and results claimed, for the model. It is now accepted that the system undergoes a form of localization with a nontrivial mobility edge. Moreover, there are strong indications for the existence of a NEE phase between the localized and ergodic regions.

Despite this, a clear picture of the actual implications of the dynamical properties of the system in connection with PT is not quite available. Our purpose in this section is to try and shed some light on this topic at least in the small-size regime, where exact diagonalization is possible. As we will see, this kind of study has little to say about the asymptotic regime, which is instead the typical playground for theoretical calculations, but on the flipside, it can orient our realistic expectations about implementing the PT algorithm in NISQ devices.

4.3.1 Model definition and previous results

We define the *random energy model* (REM) by the σ^z -diagonal Hamiltonian

$$H_{\text{REM}} = \sum_{z=0}^{2^n-1} E_z |z\rangle \langle z| \quad (4.36)$$

where the 2^n energies E_z are i.i.d. Gaussian random variables distributed according to $p(E) = \sqrt{1/\pi n} \exp(-E^2/n)$. The classical random energy model was introduced in Ref. [81] by Derrida in order to provide a simplified model that exhibits some of the glassy physics of the Sherrington–Kirkpatrick spin glass [229]. Being a collection of states with uncorrelated energies, the number of states at a given energy is given in the thermodynamic limit by

$$\rho(E) = \frac{2^n}{\sqrt{\pi n}} e^{-E^2/n}, \quad (4.37)$$

The REM is the prototypical example of a model with “golf-course” energy landscape, where the typical energy density is $\epsilon = E/n = 0$ (the flat “plateau”) while states with finite energy densities $0 < |\epsilon| \leq \sqrt{\log 2}$ are randomly scattered inside the plateau, and typically separated by an extensive distance ($d \in O(n)$) from one another. In this sense, the classical random energy model exhibits energy clustering [33,182], albeit of a degenerate kind as each “cluster” is composed of a single state. In the large- n and large-time limit, its local dynamics is exactly captured by Bouchaud’s trap model [30,54,111]. A discussion of more realistic potentials with correlated landscapes is given in Appendix D.

The *quantum random energy model* (QREM) is the transverse-field model obtained by taking the REM Hamiltonian as the problem Hamiltonian ($H_{\text{REM}} = H_p$ in Eq. (4.2)). The




Figure 4.2: Dynamical phase diagram of the QREM. The solid red line $\Gamma_{\text{MBL}}(\epsilon)$ marks the Anderson transition; the dashed red line represents the leading-order estimate $\Gamma_{\text{MBL}}(\epsilon) \approx |\epsilon|$ in the region where the approximation is at least 90% accurate. The orange line is the band edge. The purple line denotes the onset of the quantum paramagnetic region. The dot-dashed blue line represents a rough estimate of the putative non-ergodic-ergodic transition. Our analysis is carried out along the dashed black line, $|\epsilon| = 0.27$.

QREM Hamiltonian is formally equivalent to the Anderson model (2.1) with a Gaussian-distributed disordered local potential and nearest-neighbor hopping defined over the Boolean hypercube of connectivity n , where the σ^z -diagonal term H_{REM} plays the role of the random local potential and the transverse field is the kinetic term. The number of sites on this fictitious lattice is equal to $D = 2^n$. Unlike the standard approach to the Anderson model, where the hopping coefficient is kept fixed and the transition is induced by the disorder-strength parameter W , in our approach it is more useful to keep the energy scale of the local potential fixed, and induce transitions by controlling the kinetic energy parameter $\Gamma \sim 1/W$. This allows one to study the dynamics of the PT protocol through the methods of Anderson localization.

The dynamical phase diagram of the QREM has been the object of study of several previous works. In Refs. [32, 159] the mobility edge for the Anderson transition was estimated as $\Gamma_c(\epsilon) \sim |\epsilon|$ using the forward scattering approximation. This is generally believed to be correct near $\Gamma = 0$ but possibly unreliable at larger Γ as it involves a perturbative calculation of the Green's function with Γ as a small parameter. In Ref. [98] it was argued that at low enough energies and at least in an interval of Γ , the Rosenzweig–Porter model (RPM), Def. 4.2.3, should provide an effective model of the QREM. Through this mapping one finds that the QREM should exhibit all of the three dynamical phases of the RPM, *i.e.* localized, extended non-ergodic and extended ergodic (see Theorem 4.2.1). The two transitions between these phases are respectively called the *Anderson* and the *ergodic* transitions. In Ref. [235] the authors derive an estimate of the ergodic transition that coincides with the one in Ref. [98], but argue that the NEE phase is layered in an alternating sequence of two distinct subphases. The different estimates for the three phase transitions for the QREM are summarized in Fig. 4.2.

As for the population transfer protocol, following the analysis of Ref. [234] on the impurity band model, as laid out in the previous section, one would expect a polynomial speedup over random search, although a more recent work by the same authors [235] provides a precautionary statement against naively drawing that conclusion. In a different work on the quantum p -spin model [31], the PT dynamics is computed through the Schrieffer–Wolff perturbation theory in Γ . Their analysis (extrapolated to the QREM by taking the standard limit $p \rightarrow \infty$) claims that no speedup is expected on the QREM.

4.3.2 Quench dynamics in the non-ergodic phase

In this section we study the dynamics of the PT protocol described above in the following steps:

- 1) We detect the transition from a localized to an extended phase using the volume scaling of the Shannon entropy of the energy eigenfunctions in the computational basis. We observe non-ergodicity for a significant interval of Γ values.
- 2) We define a criterion for estimating the timescale t_{sat} of the PT-induced delocalization

process by observing the saturation of the Shannon entropy of the time-evolved wave function. We study the volume scaling of the saturated entropy and compare it with the scaling of entropy of the energy eigenfunctions.

- 3) We study the distribution of the saturation times t_{sat} over the disorder.

Entropy of the energy eigenstates

Multiple ways of detecting the NEE phase have been proposed in the literature (see *e.g.* Refs. [15,96,98,152,153,208]). We found that the easiest way of detecting the Anderson and the ergodic transitions numerically is the scaling analysis of the Shannon entropy of the energy eigenstates $|\psi_a\rangle$ of the QREM Hamiltonian in Eq. (4.2). Given any such eigenstate, its Shannon entropy (in the computational basis $\{|z\rangle\}_z$) is

$$S[\psi_a] = - \sum_z |\langle z|\psi_a\rangle|^2 \log |\langle z|\psi_a\rangle|^2, \quad (4.38)$$

and one defines the scaling dimension¹⁰ α_{st} through its asymptotic behavior

$$S[\psi_a] \sim \alpha_{\text{st}} \log D. \quad (4.39)$$

Since $0 \leq S \leq \log D$, the exponent α_{st} must lie between zero and one. Note that $\exp(S)$ defines the size of the “typical set” that includes the sites z that the wave function associates with higher probabilities $|\psi_z|^2$, also called the “support set” of the wave function in the physics literature [76,78]. An exponent α_{st} indicates that the number of states in the support set of the wave functions scales as $D^{\alpha_{\text{st}}}$ in the large D limit. In the localized regime the energy eigenfunctions decay exponentially away from their localization center so that most of the amplitude is concentrated in a region of size $O(D^0)$, while in the ergodic extended regime they are roughly uniformly extended over the whole system (of size $O(D)$). This means that the possible values of α_{st} can be used to identify the dynamical phases in the following way:

$$\begin{cases} \alpha_{\text{st}} = 0 & \text{localized phase} \\ 0 < \alpha_{\text{st}} < 1 & \text{non-ergodic extended phase} \\ \alpha_{\text{st}} = 1 & \text{ergodic extended phase.} \end{cases} \quad (4.40)$$

Thus, in order to assign a point (ϵ, Γ) in the phase diagram to one of the three dynamical phases, we do as follows. For each system size $n = 8, \dots, 18$ we generate a large number of disorder realizations H_J of the classical REM model of Eq. (4.36). Then, we create the

¹⁰Note that in the multifractal literature (see *e.g.* Ref. [95]), this scaling dimension α_{st} coincides with the fractal dimension α_1 , also known as the “information dimension”. However, we will make no use of this connection in the present work.

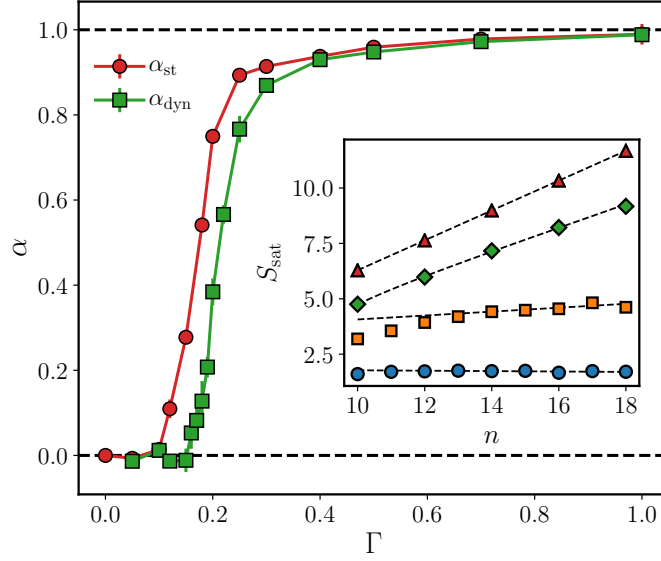


Figure 4.3: Behavior of the scaling dimension of the eigenstate entropy α_{st} (red circles) and of the dynamical entropy α_{dyn} (green squares). *Inset:* linear fit of the dynamical entropy for fixed values of the transverse field (from bottom to top, $\Gamma = 0.05, 0.16, 0.2, 0.4$).

QREM realization $H_j(\Gamma)$ by adding a transverse-field term with the given value of Γ and we use the shift-invert method to extract the energy eigenstate $|\psi_a\rangle$ of the QREM Hamiltonian whose energy is closest to $E = \epsilon n$. We compute the Shannon entropy of the eigenstate $|\psi_a\rangle$ in the σ^z basis (Eq. (4.38)) and take the average of the disorder. Thus, we obtain the typical values S_D for systems of volumes $D = 2^8, 2^9, \dots, 2^{18}$. Finally, through Eq. (4.39) one can see that a linear fit of S_D vs $\log D$ (for large enough D) will produce the α_{st} associated to the phase-diagram point (ϵ, Γ) . The results for an energy density $\epsilon = 0.27$ and $0 \leq \Gamma \leq 1$ are shown in Fig. 4.3, where one notes that α_{st} (indicated by the red circles) increases continuously from zero to one as Γ is increased from zero.

Dynamical entropy

Since Anderson's seminal paper [19], the Anderson transition was defined by two qualitatively different dynamical behaviors that can occur in a quantum system as an initially localized state is left to evolve under coherent evolution.

In this section we study numerically this delocalization process in the QREM, with an emphasis on the NEE phase. In analogy with what we did in the previous section, we will be focusing on the dynamical value of the Shannon entropy of the wave function

$|\psi(t)\rangle = \exp(-iHt) |z_0\rangle$, that is

$$S(t) = - \sum_z |\langle z|\psi(t)\rangle|^2 \log |\langle z|\psi(t)\rangle|^2. \quad (4.41)$$

At time $t = 0$ the wave function is fully concentrated on the initial state and its entropy is therefore $S = 0$. The entropy then increases as the quantum dynamics populates resonant states, up to an (eventual) stable value S_{sat} .

For our purposes, an instance is considered to have “saturated” once the distribution of the instantaneous values of the entropy becomes in a sense indistinguishable from Gaussian noise around a stable mean. More precisely, suppose that the entropy is being sampled at times $t \in \{t_1 < t_2 < \dots < t_M\}$. Now call $S[t_j : t_M] = (S(t_j), S(t_{j+1}), \dots, S(t_M))$ the sequence of entropy snapshots after time t_j for a certain instance. We say that the instance “saturates at time t_j ” if t_j is the smallest time such that the values $S(t_{j-1}), S(t_{j-2}), \dots, S(t_{j-k})$ are *all* less likely than a given probability threshold to have been sampled from a normal distribution with compatible average and variance (the arbitrary parameter $k > 1$ is used to rule out genuinely random large deviations from the mean).

In other words, for a given small $p > 0$ we compute a corresponding threshold $\xi = \sqrt{2} \operatorname{erf}^{-1}(1 - p)$, such that a standard Gaussian variable has probability p of being larger than ξ in absolute value. Then, we look for the smallest J such that

$$\phi(t_j) = \frac{|S(t_j) - \langle S[t_j : t_M] \rangle|}{\operatorname{stdev}(S[t_j : t_M])} > \xi \quad (4.42)$$

for all $j \in \{J - k, \dots, J - 1\}$, where $\langle X \rangle$ represents the arithmetic mean of process $X = (x_1, \dots, x_K)$ and $\operatorname{stdev}(X) = \sqrt{\frac{K}{K-1} \langle [(X - \langle X \rangle)^2] \rangle}$. We take $k = 4$ (3 for $n \in \{10, 11\}$), $p = 1\%$ and require that the number of samples between t_j and t_M be sufficiently large (at least 15 samples).

The above condition allows us in principle to define t_j as the saturation time. However, due to the finite sampling rate, this will result in an “aliased” distribution for the saturation time, with the loss of precision associated to constraining t_j to belong to the set of sampled times $\{t_j\}_{j=1}^M$. In order to ameliorate this effect we introduce a simple interpolation procedure.

Suppose that t_j is the smallest sampled time where the aforementioned condition holds. Then by assumption $\phi(t_j) \leq \xi < \phi(t_{j-1})$. Assuming that function ϕ is continuous, we may approximate its unsampled behavior between t_j and t_{j-1} through a linear function¹¹, and define the saturation time t_{sat} through the condition

$$\phi(t_{\text{sat}}) = \xi, \quad (4.43)$$

¹¹We actually use a linear-in-log(t) function as our sampled times are logarithmically spaced.

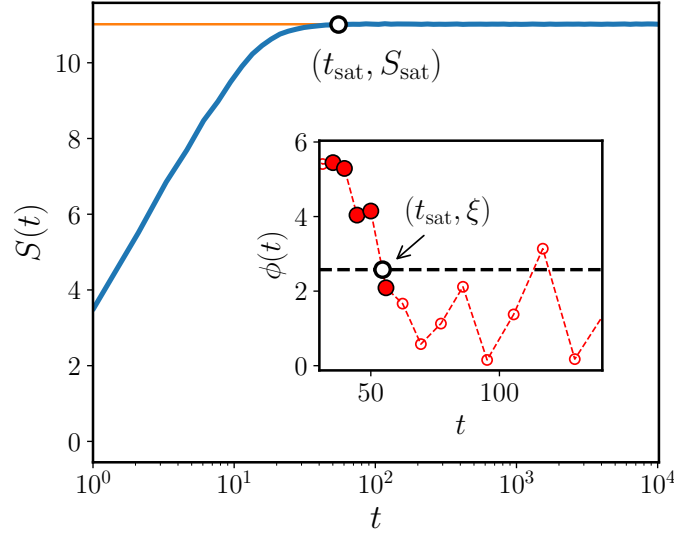


Figure 4.4: Typical behavior of the Shannon entropy $S(t)$ of the wave function (Eq. (4.41)) in the PT dynamics for a single instance with $n = 18$. S starts at zero and increases with t , up to a point of saturation $S_{\text{sat}} = S(t_{\text{sat}})$ (orange line), where it settles with small (indiscernible in the plot) random fluctuations, which we model as a Gaussian process. *Inset:* instantaneous value of the discrepancy $\phi(t)$, see Eq. (4.42), in a limited range of times. Highlighted in red are the points t_{j-4}, \dots, t_j , which define the saturation condition as explained in the text. Note the genuinely random large deviation at $t \approx 120$, which is correctly ignored by having $k > 1$.

with the solution obviously satisfying $t_{j-1} < t_{\text{sat}} \leq t_j$. This will give us a more natural distribution for the saturation time, as t_{sat} is now unconstrained.

Fig. 4.4 portrays the typical time dependence of $S(t)$ and provides a visualization of the definition of t_{sat} in terms of the function $\phi(t)$.

Fractal delocalization of the wave function

In analogy with the previous section, we use the saturation values of the Shannon entropy S_{sat} to measure the size of the region of the lattice that the wave function delocalizes over under unitary evolution. In the ergodic phase one expects that for large t , $S(t) \sim \log D$ as the wave function eventually populates all available volume, while in a strongly localized phase (*e.g.* the localized phase of the Anderson model in finite dimensions, where the RAGE theorem holds [10]), $S(t) \sim O(D^0)$ since the wave function is trapped for all times in a finite region of space surrounding the initial position. According to Ref. [234], in the NEE phase the initial state should delocalize over the common support (*i.e.* the intersection of the roughly-similar supports) of the energy eigenstates belonging to an energy miniband,

states that we previously saw have fractal supports. The intersection of fractal sets is commonly a fractal set itself [97] even though it may be characterized by a smaller scaling dimension than the one obtained from the original sets.

We define a scaling dimension α_{dyn} for the saturation values of the Shannon entropy of the wave function, in a way analogous to the static α_{st} of the previous section:

$$S_D(t_\infty) \sim \alpha_{\text{dyn}} \log D \quad (4.44)$$

where S_D is the median entropy of the time-evolved wave function for a system of volume D and we sample at a large time t_∞ such that $t_\infty \gg t_{\text{sat}}$ for every instance.

For large enough values of the transverse field, the median saturated entropy S_{sat} is nicely fit by a law of the form $S_D(t_\infty) = \alpha_{\text{dyn}} \log D + C_0 + C_{-1}/D$ (cf. [153]), where the finite-size deviation coefficient C_{-1} is always comfortably small.

However, as can be appreciated from the inset of Fig. 4.3 ($\Gamma = 0.16$, square orange markers), the small- Γ data are affected by larger finite-size effects in that they take longer to reach the asymptotic regime. These effects are not well captured by a continuous ansatz, but are better described as a sudden change of regime. Combined with the small value of the scaling dimension in that regime, this requires special care in extracting a useful estimate of α_{dyn} . We chose to restrict our fitting procedure to the $n \geq 14$ data for $0.12 \leq \Gamma \leq 0.20$, whereas for $\Gamma < 0.12$ the data is essentially too flat to detect a regime change and all sizes can be fitted. In both cases, we use a function $\alpha_{\text{dyn}} \log D + C_0$ for the fit.

With these provisions, α_{dyn} can be computed and compared to α_{st} , as shown in Fig. 4.3. The dynamical exponent has a behavior qualitatively identical to the static one, albeit shifted in the Γ axis, and it can be seen to satisfy

$$\alpha_{\text{dyn}} \leq \alpha_{\text{st}}. \quad (4.45)$$

This stands in agreement with our above intuition that the support of the time-evolved wave function essentially consists of the intersection of multiple eigenstate supports, themselves scaling in size with fractal exponent α_{st} .

Decay times for delocalization

We now study the distribution of saturation times for the PT process. Fig. 4.5a shows the histograms of the saturation times separated by system size n . Note that they exhibit a peak with a long right tail. We can use a (truncated) complementary cumulative distribution function (CCDF)

$$F(t) = \int_{t_{\text{peak}}+t}^{\infty} p(t_{\text{sat}}) dt_{\text{sat}} \quad (4.46)$$

to study the right tail of the distributions. t_{peak} is the position of the peak of $p(t_{\text{sat}})$ estimated through a Gaussian smoothing of the data. We approximate $F(t)$ using the

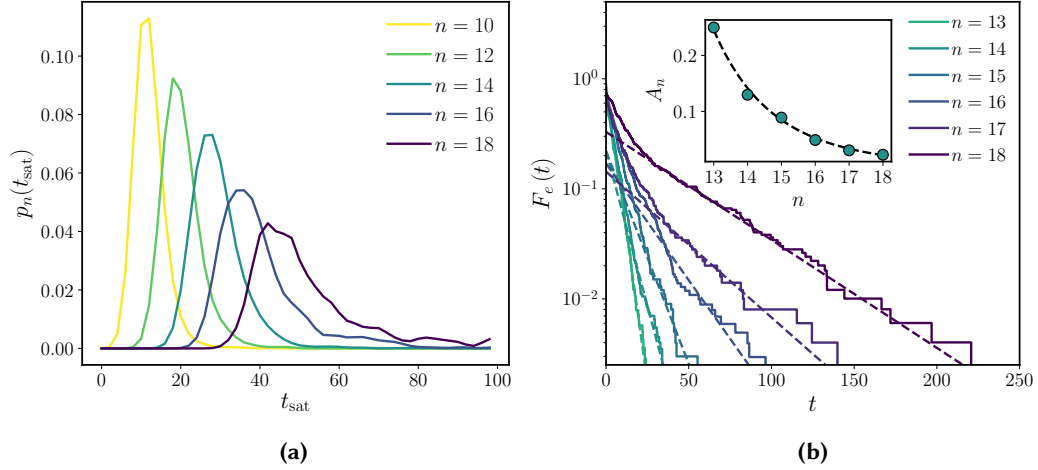


Figure 4.5: (a) Distributions of the saturation time at the NEE point ($\Gamma = 0.4$). The data was smoothed through integration against a Gaussian kernel of width $\sigma = 2$. (b) Right tails of the distributions, as described by the empirical CCDF, Eq. (4.47). Dashed lines show the exponential fits, Eq. (4.48). *Inset:* exponent A_n , as defined in Eq. (4.48), describing the decay rate of the distribution of t_{sat} at large values of its argument. The dashed line represents the power-law fit $A_n \sim an^{-b}$, with $b = 7.4(2)$.

empirical CCDF $F_e(t)$ obtained from the raw data $\{t_{\text{sat}}^{(i)}\}_{i=1}^{N_s}$

$$F_e(t) = \frac{1}{N_s} \sum_{i=1}^{N_s} \Theta(t_{\text{sat}}^{(i)} - t - t_{\text{peak}}) \quad (4.47)$$

where Θ is the Heaviside step function. We find that $F_e(t)$ is well-approximated by a two-parameter exponential functional form (Fig. 4.5b)

$$F_n(t_{\text{sat}}) \approx \exp\left(-A_n t_{\text{sat}} - B_n\right) \quad (4.48)$$

where (Fig. 4.5b, inset)

$$A_n \sim an^{-b} \quad (4.49)$$

so

$$F_n(t_{\text{sat}}) \sim \exp\left(-a \frac{t_{\text{sat}}}{n^b}\right) \quad (4.50)$$

which means that the probability density $p_n(t_{\text{sat}})$ for large values of t_{sat} is

$$p_n(t_{\text{sat}}) \sim an^{-b} \exp\left(-a \frac{t_{\text{sat}}}{n^b}\right) \quad (4.51)$$




Figure 4.6: Box plot of the distributions in Fig. 4.5a. Boxes subsume the middle half of the data. The black bar in the middle of each box is the median, with its bootstrapped error denoted by notches in the box. White squares represent the mean values. Whiskers extend from the smallest datapoint above $x_{\text{low}} = q_1 - \frac{3}{2}(q_3 - q_1)$ to the largest one below $x_{\text{high}} = q_3 + \frac{3}{2}(q_3 - q_1)$, where q_1 and q_3 are the first and third quartile, respectively. The dashed line depicts a shifted exponential fit of the median times, with $\theta_t = 0.04(2)$.

that is, even though the probability distributions p_n have thin (*i.e.* exponentially-decaying) right tails at all finite sizes n , their decay coefficients get increasingly smaller with n so the tails get fatter with increasing n .

Fig. 4.6 shows a tentative finite-size scaling of the median saturation time. An exponential ansatz with three parameters is shown (dashed line). This choice of parametrization will be motivated in Section 4.3.4.

4.3.3 PT quality: spread and spillage

In this section we start to analyze the PT dynamics as a solution-mining algorithm. In order to benchmark its performance we need to study quantitatively how the time-evolved wave function

$$|\psi_\infty\rangle = e^{-iHt_\infty} |z_0\rangle \quad (t_\infty \gg t_{\text{sat}}) \quad (4.52)$$

populates the target subspace Ω_n . In an ideal situation, one would like that 1) the wave function should be completely contained in the target subspace, and 2) it should be uniformly spread, that is

$$|\psi_\infty(z)|^2 = \begin{cases} 1/|\Omega_n| & \text{if } z \in \Omega_n \\ 0 & \text{otherwise,} \end{cases} \quad (4.53)$$

as this would provide an effective way of finding *all* target solutions efficiently by repeatedly sampling the distribution $p[\psi_\infty] = \{|\psi_\infty(z)|^2\}_z$ — a principle known as “fair sampling” [269]. This is of course an ideal limit, and is never achieved in practice. Nevertheless, we can use the displacement from this ideal case as a yardstick for PT success. We will first consider the following two quantities:

- a) the total probability $L_\Omega[\psi_\infty] = \|P_\Omega^\perp |\psi_\infty\rangle\|^2$ of finding a state outside of the target subspace (*amplitude spillage*), and
- b) the degree of uniformity of the wave function’s intensity distribution with respect to the Hamming distance, which we measure by means of a functional $U_\Omega[\psi]$ described in detail in Appendix E. This functional satisfies $U_0 \leq U_\Omega[\psi] \leq 1$ for all distributions $p[\psi]$, taking its minimal value $U_0 \approx 0.6065$ when the restriction of $p[\psi]$ to Ω is completely uniform and its maximal value 1 when it is atomic.

In Fig. 4.7 we show the behavior of these two quantities in three points of the phase diagram that we take as exemplary points for the localized ($\Gamma = 0.05$), NEE ($\Gamma = 0.4$) and ergodic behavior ($\Gamma = 1$). All points are associated to the same energy density $\epsilon = 0.27$. As a first sanity check, note that

- a) the localized case has the least amount of spillage, but is also the most inhomogeneously spread;
- b) the spread of the wave function component in the target subspace (as measured by the U_Ω functional) is essentially the same in the NEE and in the ergodic case.

In order to see why this is the case, we plot the total probability inside the target subspace resolved by the fractional Hamming distance $x = |z - z_0|_H/n$ from the initial state z_0 (Fig. 4.8); namely, the distribution

$$p_\Omega[\psi_\infty](x) = \sum_{z \in \Omega} |\psi(z)|^2 \delta_{|z - z_0|_H, nx}. \quad (4.54)$$

We can make several observations about these results. In the localized case the amplitude is progressively more concentrated on the initial state z_0 as n is increased. The PT protocol in this phase will not be able to efficiently find target states that are more than $d = xn$ spin flips away from z_0 , for large enough $0 < x < 1$. In the ergodic case, on the other hand, we see a roughly Gaussian distribution of amplitudes over the fractional Hamming distance. This is an effect of the geometry of the Boolean hypercube, where most states are approximately $n/2$ spin flips away from the initial state z_0 . The distributions seem to be essentially independent of n (within the statistical errors due to the finite sampling of the disorder). Finally, in the NEE case we see the same overall Gaussian shape of the distributions, but in this case they are larger than the ergodic case, and more importantly, they *flow favourably* (in the sense of PT) with increasing n . This is due to the fact that the amplitude is dynamically biased toward the target subspace, unlike in the ergodic case.




Figure 4.7: *Top:* distribution of the L_Ω (“spillage”) functional, discriminating ergodic regimes ($L_\Omega \approx 1$) from non-ergodic ones ($L_\Omega < 1$). *Bottom:* distribution of the U_Ω (“spread”) functional, discriminating localized regimes ($U_\Omega \approx 1$) from extended ones ($U_\Omega \approx ((1 + e^{-1/n})/2)^n$; cf. Appendix E). For an explanation of the box plot, see Fig. 4.6.

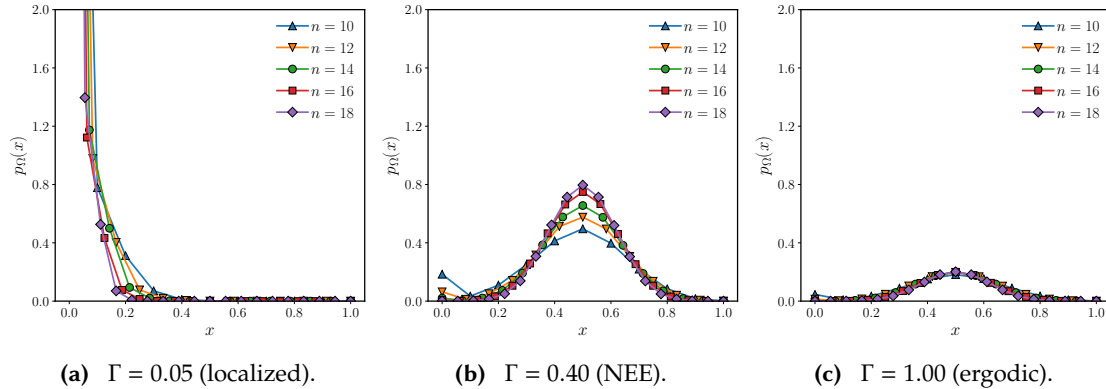


Figure 4.8: Probability distributions of finding a resonance via PT as a function of its fractional distance $x = d/n$ from the initial state, at energy density $\epsilon = 0.27$ and different values of the transverse field

4.3.4 Quantum advantage

Thus far we have learned some facts about the PT protocol in the QREM, but its application to quantum computing is practically relevant only insofar as it can be cast as a quantum algorithm that is able to outperform all classical algorithms at some specific computational task, a situation described as *quantum advantage* (or *speedup* if time is the relevant quantity). In order to show quantum advantage in the energy matching problem described in Section 4.1, one would like to compare the running time of the PT algorithm against an optimal classical algorithm designed to find all states at a given energy. This is generally a significantly understudied problem. Some very specific models have provably efficient algorithms that sample the Gibbs distribution for a given temperature (e.g. the planar Ising spin glass [249]), so that if one is able to compute the Legendre transform between energy E and temperature T then in the large n limit one could obtain all states at a given energy by sampling the Gibbs state of the associated temperature T_E . Failing that, one is reduced to using a general-purpose stochastic local search algorithm such as simulated annealing (cf. Section 1.3.2).

Unfortunately, such comparison is not very useful in the case of the REM due to its trap-model dynamics: its local dynamics is made up of a sequence of thermally-activated events where the system is excited out of a low-energy state to the flat plateau (the classical states with $\epsilon = 0$), followed by a random walk on the plateau until it falls again into another (random) low-energy state with a random energy E . If E lies inside the target energy shell, then we have found a target state of the energy matching problem and we are done. If not, we need to wait for the system to be thermally excited out of this state and the search process starts anew. Even assuming one can speed up the excitation events so that they only take $O(1)$ time, there seems to be no way to avoid having to perform a random walk on

the plateau in order to find a new low-energy state. Due to the flatness of the plateau and the random relative positions between states of non-typical energy density, this stochastic local search can ultimately be seen as a *global* random search.

In view of these considerations, we adopt the following benchmark metric. For a given realization of the REM Hamiltonian H_p , we define two “oracles”, \mathcal{O}_{qu} and \mathcal{O}_{cl} , that take as input an initial bit string z_0 and attempt to produce a target string $z' \in \Omega_n$. The first is a quantum oracle that applies a PT evolution to $|z_0\rangle$ for a time $t = t_{\text{sat}}$, with some n -independent driver strength Γ , and then measures the final state $|\psi(t)\rangle$ in the computational basis. The second oracle is a classical procedure that discards the initial string z_0 and simply samples a new one $z \in \{0, 1\}^n$ at random, with uniform probability $p(z) = 1/2^n$. The probability of success for *one call* of each oracle is

$$P_{\text{qu}} = \sum_{z \in \Omega_n} |\langle z | \psi(t) \rangle|^2, \quad (4.55)$$

$$P_{\text{cl}} = \frac{|\Omega_n|}{2^n}. \quad (4.56)$$

We define the *gain*

$$G = \frac{P_{\text{qu}}}{P_{\text{cl}}} \quad (4.57)$$

as the ratio of these two probabilities. The gain is a random variable due its dependence on both the problem instance H_p (which determines the set Ω_n) and the choice of initial string z_0 (which determines $|\psi(t)\rangle$).

For each size $n = 8, \dots, 20$ we sampled a number of disorder realizations of the REM Hamiltonian H_p and for each of them we selected the string z_0 with an associated energy closest to $E = 0.27n$. The probability P_{cl} can be computed by simple inspection of the random energies in the realization H_p , while to compute P_{qu} we simulated the PT protocol numerically. We performed a finite-size scaling of the data thus obtained (Fig. 4.9) using the ansatz

$$G(n, \Gamma) = (Ae^{\alpha n} + B)\tilde{G}(\tilde{\Gamma}(n)) \quad (4.58)$$

where the effective $\tilde{\Gamma}(n)$ is given by

$$\tilde{\Gamma}(n) = \Gamma - \Gamma_\infty + \frac{C}{n^{1/\nu}}. \quad (4.59)$$

Given the available data, fitting three parameters results in a poor estimate for the uncertainties on C and ν . Therefore, we prefer to fix one of them while keeping the other one free. Observing that the gain peaks approximately in correspondence of the Anderson transition, we chose to utilize the same value of ν describing the finite-size shift of the mobility edge, which is known to be $0.3 \leq \nu \leq 0.5$ [32]. We use $\nu = 0.4$ so that $1/\nu = 2.5$ ¹².

¹²We remark that even if we leave ν as a free parameter in the fit, we obtain a value of approximately $1/\nu = 2.53$.




Figure 4.9: Curve collapse of the median gain for system sizes $n = 14$ through 20.

We obtain the following values for the fitting parameters:

α	0.084(5)
A	5.3(8)
B	$-0.27(13) \times 10^1$
C	$2.58(12) \times 10^1$
Γ_∞	0.242(1)

This means that

- the median gain peaks at a size-dependent value of $\Gamma_{\text{peak}}(n)$ that for $n \rightarrow \infty$ is in quite good agreement with the estimate of the Anderson transition's mobility edge ($\Gamma_c \approx 0.27$ following the forward scattering approximation). The thermodynamic-limit position of the peak Γ_∞ is shifted by a finite-size effect $\Delta\Gamma_\infty(n) = -Cn^{-1/\nu}$. This shift was already observed in Ref. [32] in the finite-size scaling analysis of the mobility edge of the QREM using the forward-scattering approximation, as well as other quantum phase transitions in disordered transverse-field models [189].
- the height of the peak scales like $G_{\text{peak}} \sim e^{\alpha n}$ with $\alpha > 0$. This is a query-complexity asymptotic advantage of the quantum PT oracle over the random search oracle.




Figure 4.10: Peak median gain for $n = 13, 14, \dots, 20$, fitted with a shifted exponential $Ae^{\alpha n} + B$ with parameters $A = 5.3(8)$, $\alpha = 0.084(5)$, $B = -27(13) \times 10^{-1}$. Inset: position of the peak median gain for $n = 14, 16, 18, 20$. The scaling Ansatz Eq. (4.59) suggests a value $\Gamma_{\text{peak}} \approx 0.24$ in the thermodynamic limit.

Non-oracular advantage

In order to assess the absolute (*i.e.* non-oracular) advantage of the PT protocol over stochastic random search one has to compute the full runtime of both algorithms by multiplying the number of oracle calls by the time required to implement a single oracle call. Equivalently, one can study the quantity

$$K = \frac{P_{\text{qu}}}{t_{\text{qu}}} \frac{t_{\text{cl}}}{P_{\text{cl}}} = G \frac{t_{\text{cl}}}{t_{\text{qu}}}, \quad (4.60)$$

where $t_{\text{qu}}, t_{\text{cl}}$ are the times needed to implement one call to the quantum and classical oracles, respectively. For the classical oracle, a random search requires a time $t_{\text{cl}} \in O(n)$ as one only needs to generate n random bits (linear time for a classical probabilistic Turing machine). For the quantum oracle, t_{qu} is the saturation time t_{sat} we studied in Section 4.3.2. Therefore $K(n) = knG(n)/t_{\text{sat}}(n)$ with k an $O(n^0)$ prefactor. We consider a setup where, for fixed $\Gamma = 0.4$ (that represents a good choice of parameter setting without being optimally tuned) we evolve the initial state according to the PT protocol for a time $t = t_{\text{sat}}(n)$ and then we sample the wave function in the computational basis. Quantum advantage is decided by the competition of the gain G and the saturation time t_{sat} .

Unfortunately, the data at finite sizes $n = 10 - 20$ show only a weak dependence on n so that the functional form of these two quantities is hard to extract conclusively. On

physical grounds, we expect an exponential scaling of both G and t_{sat} , so we fit through functions of the form

$$f(n) = ae^{\theta n} + b, \quad a, b, \theta \in \mathbb{R}. \quad (4.61)$$

The values of the θ exponents control the rate of divergence of the leading terms of these quantities in the limit $n \rightarrow \infty$. We use the notation θ_t and θ_g for the time and gain exponent, respectively. We obtain the values

$$\begin{aligned} \theta_g &= 0.09(4), \\ \theta_t &= 0.04(2). \end{aligned}$$

Thus we have the following approximate scaling

$$K = k \frac{n G(n)}{t_{\text{sat}}(n)} \sim \exp[(\theta_g - \theta_t)n] \approx \exp[0.05(4)n]. \quad (4.62)$$

While this would imply an asymptotic speedup in the $n \rightarrow \infty$, the prefactor of n at the exponent is very small (in fact, statistically compatible with zero) and we believe that the most reasonable interpretation of these results is that the PT protocol and random search scale approximately in the same way, at least at the sizes we have been able to study.

Indeed, for such a small value of $\delta\theta = \theta_g - \theta_t$, the linear factor kn from t_{cl} will affect the behavior of K on a significant range of pre-asymptotic system sizes, $n \lesssim 1/\delta\theta$, regardless of its true asymptotic scaling. However, this does not seem to make much of a difference for the system sizes we can observe: if we plot the median values of the K value directly (instead of fitting the exponential scaling of G and t_{sat} separately as in Eq. (4.62)), then the corrections to the exponential terms coming from the parameters a, b in Eq. (4.61) and the linear term $t_{\text{cl}} = kn$ still cancel out any significant difference between G and t_{sat} , and the slope of K turns out to be extremely small (see Fig. 4.11)

Our conclusion is that in the QREM, the PT protocol and random search scale approximately in the same way for problem sizes $n \leq 20$. The small curvature of both the time and gain data suggests that going to moderately larger sizes, as in some of the largest numerical analysis currently available in the many-body localization literature [206], is unlikely to improve our analysis by a significant degree. Any eventual asymptotic advantage would be detectable only at unachievably larger sizes.

4.4 Conclusions

In this chapter, the application of quantum computing to a useful computational primitive called energy matching, as defined in Section 4.1, was surveyed in detail.

In Section 4.2 we have reviewed one of the first and most thorough theoretical analyses of the population transfer (PT) algorithm, which considers energy matching in a toy “impurity band” Hamiltonian [234]. We have described how a semiclassical approximation,




Figure 4.11: Distributions of the value of K , as defined in Eq. (4.60). For an explanation of the box plot, see Fig. 4.6.

perturbative in the IB width as well as the driver strength, leads one to consider a random ensemble of “preferred-basis Lévy matrices”, a fat-tailed analogue of the Rosenzweig–Porter model. A Fano–Anderson treatment of the PT process leads to an estimate of the decay rate, which is itself a fat-tailed random variable whose distribution can be estimated by means of a self-consistent nonlinear equation whose solution can be obtained via the cavity method. The resulting estimate for the typical PT time indicates a quasi-Grover speedup, *i.e.* quadratic over classical random search up to an anomalous exponent that is vanishing as $\sim 1/\Gamma^2$ in the transverse field.

In Section 4.3 we have studied numerically (up to system sizes of $n = 20$ quantum spins) the PT dynamics in the quantum random energy model [203]. We summarize our methods, findings and conclusions hereunder.

For a horizontal line in the (ϵ, Γ) phase diagram of the model, corresponding to fixed energy density $\epsilon = 0.27$, we computed the scaling of the support size of both the energy eigenfunctions, and the wave functions obtained at the end of the delocalization process induced by the system’s coherent dynamics applied to a localized initial state. For increasing values of Γ , we observed in both cases a continuous crossover of the scaling dimension α from a localized ($\alpha = 0$) toward an ergodic regime ($\alpha = 1$), which indicates the presence of a non-ergodic extended phase.

We further studied the timescale t_{sat} required for the PT delocalization process to complete, as a function of the system size n . Over the disorder, this time follows a unimodal distribution with a long right tail that gets fatter with increasing n . The median

value of t_{sat} shows an exponential dependence of on the system size, $t_{\text{sat}} \sim \exp(\theta n)$.

We computed the probabilities $P_{\text{qu}}, P_{\text{cl}}$ of obtaining a state in the target microcanonical shell (containing the states with the same energy density as the initial state), using respectively the saturated PT dynamics for a given value of Γ , and random search. The ratio $G = P_{\text{qu}}/P_{\text{cl}}$ (“gain”) of these probabilities corresponds to a query-complexity comparison between a single use of a black box that performs the PT protocol to completion, and then samples the resulting wave function in the σ^z basis, and a single use of a black box that samples classical states uniformly at random. We observe that the gain grows with n in the regime of sizes that is accessible to our numerics for all values of Γ we considered — *i.e.* the quantum PT oracle increasingly outperforms random search — with a peak that appears in proximity of the critical point of the Anderson transition.

However, if we compare the actual (*i.e.* non-oracular) runtime of the two algorithms by dividing the probability of success of the black boxes by the time required to implement one call to the black box, we find that the scaling of the gain $P_{\text{qu}}/P_{\text{cl}}$ and the scaling of the inverse times ratio $t_{\text{cl}}/t_{\text{qu}}$ roughly cancel out: at the system sizes considered in this work, PT and random search seem to scale with n in approximately in the same way.

We believe that the final considerations we can extract out of the above results are twofold, one concerning the choice of parameters for the PT protocol and another concerning the overall efficiency of the PT protocol as a quantum algorithm for energy matching. For what concerns the parameter setting, the fact that the optimal choice of Γ seems to coincide with the critical point of the localization–delocalization transition of the QREM suggests that the parameter setting problem for Γ can be reduced to finding the mobility edge for the Anderson transition. This is a problem that has undergone extensive studies in the literature on many-body localization. Our conjecture is that this optimality result should generically hold true for any other combinatorial optimization problem to which PT can be applied. For the second PT parameter, the final evolution time t_{fin} , the situation is less clear. The choice of an optimal time is hard to do on a case-by-case basis, as we found no obvious property of a REM instance that correlates with its saturation time t_{sat} . At this point we are only able to suggest that exploratory runs should be made using randomly generated instances in order to estimate a fixed percentile of the distribution of t_{sat} (say, its median), so that one is guaranteed that for that choice of t_{fin} , a finite fraction of instances will have reached saturation and the optimal performance of PT protocol will be reached with finite probability. Clearly, more work will be necessary in order to elucidate this point before PT can be ready for practical applications.

The second point concerns the efficiency of PT for energy matching: according to our findings, PT does not show significant quantum advantage on the QREM at the sizes that are likely to be accessible to near-term quantum technologies. This observation is true of the entire NEE phase. This leaves the door open for a few potential scenarios that we are unable to resolve at the moment:

- 1) PT never exhibits asymptotic advantage in the QREM, for any energy density what-

soever. This would confirm the asymptotic results obtained by the authors of Ref. [31] by keeping the leading-order term of the Schrieffer–Wolff perturbation theory in the small parameter Γ .

- 2) PT shows asymptotic advantage, but only for a restricted interval of energy densities in the NEE phase. In general, one would expect these energies to be neither too close to the edges of the spectrum nor to the center. Currently, there is not to our knowledge a way of detecting these “good” energies from finite-size data (all approaches consider asymptotic studies in n), except for exhaustive search.
- 3) PT has a (perhaps even “quasi-Grover”, as in the QIBM) quantum advantage in the QREM, but the system sizes we can access are simply too pre-asymptotic: the scaling we are seeing at these system sizes is not representative of the one emerging in the $n \rightarrow \infty$ limit.

We believe that this question will not be resolved at the sizes $n = 20 - 30$ that seem to be the current limit of numerical methods, or will likely be achieved in the next few years. In order to decide among the scenarios above it seems reasonable to us that better numerical methods need to be developed in order to approximate the effective coherent dynamics of the QREM in a fixed energy shell so that much larger system sizes can be investigated. Downfolding techniques like the ones developed in Ref. [234] are likely to prove useful, when paired with a robust way of estimating the error of the effective propagator for fixed time t and finite size n .

One interesting line of future research involves the experimental realization of the PT protocol on quantum hardware. While unfeasible in our case due to the nonlocal nature of the REM Hamiltonian, this is in principle achievable on other finite-connectivity combinatorial models such as those discussed in Appendix D. Indeed, beyond the QREM, we believe it would be interesting to study the PT dynamics in models with correlated disorder. In these cases the expectation is that the scaling of the gain will be worse compared to the QREM, as the comparison must be made with classical algorithms or heuristics that can exploit the correlated energy landscape and therefore outperform random search. However, the inverse time ratio $t_{\text{cl}}/t_{\text{qu}}$ in Eq. (4.60) is likely to scale more favorably: the runtime t_{cl} for these probabilistic algorithms to produce one candidate solution for the energy matching problem will likely grow exponentially with the system size — as is expected to happen for essentially any “hard” optimization problem — unlike what happens in random search, where one such attempt only requires the generation of n random bits, $t_{\text{cl}} = kn$. These local models are particularly attractive as they can be simulated directly on quantum hardware for much larger system sizes than the ones accessible to numeric methods, possibly solving the question of the asymptotic behavior of PT. Digital quantum computers are limited by the fact that they can execute programs only in the form of quantum circuits: the simulation of continuous-time processes on digital quantum hardware almost always requires the discretization of the real-time unitary propagator *e.g.* via a

Trotter decomposition (cf. Section 3.2.1), even in case of time-independent Hamiltonians such as the ones used in the PT protocol. This significantly limits the final time that can be reached by the simulation, as the elimination of Trotter errors necessitates the use of quantum circuits of non-negligible depth. Analog quantum computing technology, on the other hand, natively implements continuous-time quantum evolution, and at this point seems to be tantalizingly close to being able to simulate the PT protocol. However, currently available quantum annealers such as the D-Wave machine were not designed with PT in mind. Consequently, their annealing schedule — while more complicated than PT — is currently too rigid and cannot accommodate long intervals where the transverse field is kept constant, that are instead a necessary element of PT. We believe that the development of high-coherence analog quantum computing hardware with a broader range of applicability will greatly benefit the study of population transfer protocols for energy matching.

Finally, there remains the unexplored option of using the PT dynamics for a hybrid “PT-assisted stochastic search”, where quantum PT moves are alternated with classical stochastic search (*e.g.* some form of simulated annealing) either for the purpose of energy matching or for combinatorial optimization, as mentioned *e.g.* in Ref. [142]. This is an interesting approach that might boost the performance of these classical algorithms in specific conditions.

Closing remarks

Several and variegated topics were presented in the scope of this thesis, starting from a general introduction to ergodic theory in Chapter 1, including a discussion of blatantly non-ergodic systems such as spin glasses and their link to the theory of optimization, proceeding in Chapter 2 to an overview of how quantum systems localize as a consequence of disorder, and how symmetry plays a crucial role in this localization phenomenon, and finally discussing in Chapters 3 and 4 how stochastic and quantum dynamics can be used to our advantage in the solution of optimization problems.

The underlying motif of this work has been the violation of the ergodic hypothesis which rests at the basis of equilibrium statistical mechanics. Despite the great beauty and dignity of the theory, none of the fascinating topics that were investigated in these pages would have been possible to appreciate without venturing outside the safe haven of equilibrium. A famous brocard attributed to Niels Bohr states that the opposite of a fact is a falsehood, but the opposite of a profound truth may well be another profound truth. It is a testament to the theory that removing the cornerstone of its edifice does not provoke its collapse, but rather uncovers a whole new world of even more remarkable character. What we analyzed or just mentioned throughout these chapters offers a glimpse into this novel world of out-of-equilibrium physics.

In these closing remarks, I will take a final look at the main topics of interest that were touched on during my studies. As the conclusions specific to each work were already presented at the end of the respective chapters, I will instead now maintain a more broad-brush view and comment with greater liberty on the general context.

Let us first consider Chapter 2, where the interplay between disorder and symmetry was considered. Our current understanding of generic ergodicity violation in interacting quantum systems (that is, excluding the rather separate realm of Bethe-integrable systems) is essentially limited to many-body localization, which is in turn recapitulated in the theory of quasilocal integrals of motion. This theory successfully predicts the most salient characteristics of MBL systems, including the entanglement properties of eigenstates and time-evolved product states as well as the lack of transport of local charges.

The theory is however incomplete. For instance, the LIOM description assumes that the entirety of the spectrum is localized, and is unable to offer a good explanation of the delocalization transition and the appearance of mobility edges. Currently, our understanding of the MBL–ergodic transition is limited to phenomenological models attempting to capture what are deemed the most relevant features of the transition, *e.g.* the appearance of Griffiths regions in approaching the MBL phase. With the apparent inadequacy of perturbation theory, numerical and renormalization group methods have been attempted, but a proper understanding has not been reached yet. A related question concerns the transport properties in the ergodic side of the transition, with many results suggesting the existence of broad subdiffusive regimes. In this regard as well, an explanatory theory is still missing.

The problem of subdiffusion, and in general the characterization of the dynamics, is therefore a very interesting direction to pursue in our study of isotropic disordered models. The analysis of Section 2.4, albeit thorough, was limited to spectral properties, which in a sense concern the infinite-time behavior of the system. As stressed multiple times in Chapter 1, this does not tell the whole story for a non-ergodic system, whose approach to the diagonal ensemble is pathologically slow. Studying the finite-time behavior of simple initial states would surely provide new insights into the problem, although one must keep in mind the usual limitations of numerical simulations, which are by force constrained to relatively short times and small sizes. That said, features such as the post-quench evolution of the entanglement entropy and local observables should be usefully characterizable even through exact diagonalization methods, especially with the efficient Hilbert-space reduction and massive parallelization boasted by our code base. In addition, the low entanglement present in the system at high disorder makes it potentially amenable to DMRG techniques (although these cannot easily be adapted to nonlocal bases such as the tree basis). As mentioned in Section 2.4.5, other non-Abelian symmetries such as $SU(N \geq 3)$ could be considered as well. This would require a new code base and classifies therefore as a longer-term goal.

Ultimately, the importance of a study building on these ideas lies in its potential to uncover a whole new class of non-ergodic systems, including (but not necessarily limited to) disordered systems with non-Abelian symmetries. As we found in our analysis, the crucial feature of this kind of systems appears to be the approximate *tree tensor structure* of their eigenstates, which is enforced by the representation theory of the symmetry group. Namely, trees are generated by the application of fusion rules between representation spaces — a purely algebraic diktat. Eigenstates are constrained to be linear combinations of tree states, and in the strongly-disordered limit, as we argued from SDRG considerations, these linear combinations are approximately trivial (*i.e.* tree states are already almost eigenstates of the Hamiltonian). Since trees have logarithmic-in- L entanglement, the area law of MBL must be violated as a consequence and disorder, while still impeding ergodicity, is unable to fully localize the system, resulting in a different kind of non-ergodic phase.

The fusion of representation spaces is not a special feature of $SU(2)$, but can be ex-

tended to any semisimple Lie algebras. In general, the Clebsch–Gordan coefficients will be replaced by different numbers, and the partial fusions will be labeled by several Dynkin labels instead of just the spin, but this does not spoil the reasoning. In fact, one might conceive of cases where tree states arise for reasons other than the representation of a symmetry group. An example could be a model of non-Abelian anyons, which has no symmetry group but rather a *quantum group*. The question is open for the posterity.

Let us move on to Chapters 3 and 4. The topic of quantum supremacy, with its promise to bring about a new technological revolution of tremendous impact, is of central interest in today’s industrial research. Big competitors such as Microsoft, IBM, Intel, Google and Amazon have all been reserving generous budgets for research and development in quantum computing, and in some cases opened the arena to the larger public via cloud services for remotely accessing their exotic and otherwise unaffordable hardware. With the theory barely one or two scores young, and the practice even younger, the expectations for quantum computing are much brighter than its present-day achievements, despite existing disagreements on realistic time frames for their actualization. Currently, a lot of effort is being devoted to the all-important task of creating universal standards for the practicing “quantum computerist”, as well as developing the equally fundamental theory of quantum compilation. On the practical side, the first assembly specifications and software development kits for quantum operating systems have been published already.

In this fast-paced context of excitement and discovery, optimism for the future remains a fundamental ingredient, at least until convincing evidence is presented that quantum computers will indeed scale quickly and robustly enough to prove themselves useful in real-world applications. As it stands, classical algorithms possess a substantial advantage over quantum ones due to their head start of several decades, and have so far been quick to quell excesses of optimism whenever a quantum speedup was alleged, even in rather *ad hoc* tasks designed to their detriment. Even neglecting highly specialized algorithms, general-purpose metaheuristics such as simulated annealing appear to remain competitive despite the looming presence of quantum annealers currently marketed or in the making. Likewise, quantum routines such as population transfer have yet to be convincingly performed even in proof-of-concept setups, much less shown to outperform classical protocols or at least augment them (in the guise of “quantum subroutines” for hybrid classical/quantum algorithms).

One may argue that at least in some cases, such as Grover’s algorithms, or Shor’s algorithm if $P \neq NP$, a speedup exists *provably*. But even in the presence of a rock-solid asymptotic speedup, a *practical* speedup is only achieved when the problem size is large enough for the asymptotic term to take over. In the current regime, prefactors are extremely important and may be as large as $\sim 10^{11}$ in favor of classical computers, as Neven pointed out recently.¹³ It is therefore a moot point, as of today, to argue from asymptotic scalings. One of the only tasks where current-day quantum computers seem to be excelling is

¹³During the Google Quantum Summer Symposium 2020.

random number generation, something which classical computation should be expected to have a hard time replicating due to the fundamentally “true random” character of the quantum theory, which has no analogue in classical mechanics.

In conclusion, while justifiably sharing a good part of the optimism for the prospects of quantum computation in the near future, we must balance it with a healthy dose of cautiousness. We have tried to do so in this thesis by recognizing the limitations of quantum algorithms as well as the merits of their classical competitors. This “apophatic” approach to near-term quantum computation is less easily heralded in headlines and grant proposals, but makes just as important a part of the research effort as its opposite, if not more so. That said, what we have found serves not as a roadblock, but as a reminder to keep our expectations realistic. While diffusion Monte Carlo does display huge potential in tackling certain “rugged landscape” problems, it has to contend with hard complexity barriers by requiring very large walker populations for accurate predictions, generally exponentially so in the system size. Furthermore, DMC is incapable of dealing with non-stoquastic Hamiltonians — a limitation not shared by quantum annealers. It is not hard to envision future quantum machines leaving DMC and other stochastic algorithms far behind in certain classes of problems. Likewise, although population transfer may not live up to expectations in the random energy model, we have good reason to be optimistic about correlated models, where, as explained in Section 4.4, the oracular advantage of PT may be harder to strip out than in REM.

Acknowledgments

I want to extend my full and most heartfelt gratitude to my family for being by my side throughout the years of my doctorate, which, for all their irreplaceable value, were not always the kindest on my bodily health. Their steadfast presence and unwavering support, I could not have done without.

A similar role I must recognize in my closest friends, especially Francesco, Luca and Luca, whose loyalty and camaraderie I could always rely on, and who have supplied to me no shortage of enjoyment in times good and bad despite the frequent geographical divide. I also want to acknowledge Andrea and the great discussions we had in front of a beer.

Many a doctoral student is often oscillating between zeal and acedia toward the subject of his studies, and I found myself no different. I want to acknowledge my friends and colleagues at SISSA and ICTP, especially Giuliano, for keeping me oriented and motivated with stimulating discussions and, most importantly, through the example of their passion. I would also like to thank my professors, especially Andrea Gambassi, Pasquale Calabrese and my supervisor Antonello Scardicchio, for their guidance, as well as my referees Marko Žnidarič and Chris Laumann.

Part of my research work was performed during a six-month internship at the NASA Ames Research Center in Mountain View, CA. It was a wonderful experience and I am profoundly obliged to USRA for sponsoring my internship. Special thanks go to Saba and Zeal for their assistance during the whole period. I am equally grateful to a lot of people alongside whom I have had the privilege of working, in particular my co-interns Sathya, Ryan, Andrea and Michael and my good friend and mentor Gianni. I would be remiss not to thank Eleanor, Davide and Salvatore as well, whose experience in their respective fields is only matched by their kindness.

I am also in great debt to Alex and his family, who went far beyond the call of duty to make me feel welcome in Santa Clara County, and whom I am humbled to now call good friends. They, and so many more in the St. Joseph parish have shown me incredible warmth and made my six months in California all the more meaningful.

Finally, but most importantly, I give thanks to the Lord our God for the majesty of His creation, whose surface I have but started to scratch in my studies. *Cuius nomen sit benedictum ex hoc nunc et usque in sæculum. Amen.*

The Jordan–Wigner transformation

The Jordan–Wigner transformation [138] is a standard tool in statistical physics that converts a one-dimensional fermionic model into a spin model, or vice-versa. This allows physicists to map different models into each other and thereby amplifies the import and scope of analytical and numerical methods which may only be available for one of the two cases. For example, fermions are notoriously harder to implement in computer simulations due to the lack of a tensor product structure in their matrix representations, so it is convenient to transform them into spins beforehand, whose representation space is easily constructed via ordinary Kronecker products.

Consider the fermionic algebra generated by the annihilation operators $\{c_i\}_{i=0}^{L-1}$, acting on the local basis as

$$c_i |0_i\rangle = 0, \quad c_i |1_i\rangle = |0_i\rangle \quad (\text{A.1})$$

and obeying the canonical anticommutation relations

$$\{c_i, c_j\} = 0, \quad \{c_i, c_j^\dagger\} = \delta_{ij}, \quad (\text{A.2})$$

where $\{a, b\} = ab + ba$. We call $n_i = c_i^\dagger c_i = n_i^\dagger$ the i -th number operator.

From the elementary property

$$[c_i, n_j] = c_i c_j^\dagger c_j - c_j^\dagger c_j c_i = c_i c_j^\dagger c_j + c_j^\dagger c_i c_j = \{c_i, c_j^\dagger\} c_j = \delta_{ij} c_i \quad (\text{A.3})$$

and its conjugate

$$[c_i^\dagger, n_j] = -[c_i, n_j]^\dagger = -\delta_{ij} c_i^\dagger \quad (\text{A.4})$$

we see that all numbers operators commute with each other:

$$[n_i, n_j] = c_i^\dagger [c_i, n_j] + [c_i^\dagger, n_j] c_i = 0. \quad (\text{A.5})$$

We define the *Jordan–Wigner transformation*

$$s_i^+ = e^{i\phi_i} c_i \quad (\text{A.6a})$$

$$s_i^z = \frac{1}{2} - n_i, \quad (\text{A.6b})$$

where

$$\phi_i = \pi \sum_{l=1}^{i-1} n_l, \quad (\text{A.7})$$

and demonstrate that the s_i^a operators define a Pauli algebra:

$$[s_i^z, s_j^\pm] = \pm \delta_{ij} s_i^\pm, \quad [s_i^+, s_j^-] = 2\delta_{ij} s_i^z, \quad (\text{A.8})$$

where $s_i^- = (s_i^+)^\dagger$.

Notice that for the first identity we only have to consider one sign, as the other follows trivially from conjugating the equation. We have

$$[s_i^z, s_j^+] = [-n_i, e^{i\phi_j} c_j] = e^{i\phi_j} [c_j, n_i] = \delta_{ij} e^{i\phi_j} c_i = \delta_{ij} s_i^+, \quad (\text{A.9})$$

where we have used the number commutativity Eq. (A.5) to take $e^{i\phi_j}$ out of the brackets. As for the other commutator, we first note the property

$$e^{\pm i\pi n_l} = 1 - 2n_l, \quad (\text{A.10})$$

which can be verified by applying both sides to the basis states. Combined with Eq. (A.3), this gives

$$e^{-i\phi_j} c_i = \begin{cases} c_i e^{-i\phi_j} & \text{if } i \geq j \\ -c_i e^{-i\phi_j} & \text{if } i < j. \end{cases} \quad (\text{A.11})$$

The first line is obvious from the fact that ϕ_j does not contain i , so the whole exponential commutes with c_i . For the second line, one needs to notice that $c_i = -c_i(1 - 2n_i)$, as can be checked directly from the properties of the fermionic operators. Clearly, an analogous result holds for $c_j^\dagger e^{i\phi_j}$ by conjugation of the above.

The calculation of the commutator is now very simple:

$$\begin{aligned} [s_i^+, s_j^-] &= e^{i\phi_i} c_i c_j^\dagger e^{-i\phi_j} - c_j^\dagger e^{-i\phi_j} e^{i\phi_i} c_i \\ &= e^{i\phi_i} c_i c_j^\dagger e^{-i\phi_j} - c_j^\dagger e^{i\phi_i} e^{-i\phi_j} c_i \\ &= \begin{cases} e^{i\phi_i} \{c_i, c_j^\dagger\} e^{-i\phi_j} = 0 & \text{if } i \neq j \\ c_i c_i^\dagger - c_i^\dagger c_i = 1 - 2n_i = 2s_i^z & \text{if } i = j \end{cases} \\ &= 2\delta_{ij} s_i^z, \end{aligned} \quad (\text{A.12})$$

where in the $i \neq j$ case I used that if $i \neq j$ then one among $\{e^{i\phi_i}, e^{-i\phi_j}\}$ commutes with $\{c_j^\dagger, c_i\}$, respectively, while the other anticommutes, producing an overall minus sign.

This proves that the operators s_i^\pm and s_i^z , as defined in Eq. (A.6), are indeed a family of spin operators commuting like Pauli operators on the same site and trivially on different sites. In view of Eq. (A.6b) the local number-basis states $\{|0_i\rangle, |1_i\rangle\}$ can be interpreted after the mapping as the spin states $\{|\uparrow_i\rangle, |\downarrow_i\rangle\}$, respectively. In the new basis, the spin operators assume (locally) their habitual Pauli form,

$$s^x = \frac{1}{2} \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad s^y = \frac{1}{2} \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \quad s^z = \frac{1}{2} \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix},$$

$$s^+ = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}, \quad s^- = \begin{pmatrix} 0 & 0 \\ 1 & 0 \end{pmatrix}.$$

Although from glancing at the local states one may think that nothing has been accomplished other than a simple relabeling of the basis, the effect of the transformation is immediately clarified by considering many-particle states. Recall that the basis we started with had a fermionic character. In particular, the appropriate anticommutation relations between particles could not be represented in terms of simple Kronecker products between elementary matrices, which is instead the case for $\{s_i^a\}_{i,a}$.

This conversion between anticommuting and commuting behavior was accomplished by attaching to the ladder operators c_i and c_i^\dagger an appropriate *string operator*

$$e^{i\phi_i} = e^{in_0} e^{in_1} \dots e^{in_{i-1}}, \quad (\text{A.13})$$

whose action can be interpreted as correcting the sign discrepancies that arise as a consequence of swapping consecutive particles instead of anticommuting them. The string operator is, of course, highly nonlocal, having one end attached to the edge of the chain and the other deep in the bulk, and therefore the Jordan–Wigner transformation in general turns a local theory of fermions into a nonlocal theory of spins. However, due to the involutory property of the building blocks e^{in_0} , string operators have the very fortunate feature of combining into *local* phases when the Hamiltonian is one-dimensional and nearest-neighbor (with the possible exception of edge terms depending on the chosen boundary conditions).

Unfortunately, the presence of string operators does mean that models in higher dimensions cannot in general be Jordan–Wigner transformed without spoiling the locality of the Hamiltonian. That said, the transformation may still be viable in some simple non-tree geometries: in Section 2.3.1, for instance, a two-fermion chain (which can be seen as a single-fermion model on a ladder geometry) is successfully transformed by accurately choosing the definition of the string operators.

As a final note, we exhibit the inverse transformation. Using Eq. (A.10) and the fact that $e^{i\phi_i}$ factorizes due to Eq. (A.5), we can express the string operator in terms of spins and directly invert Eq. (A.6):

$$c_i = \left(\prod_{l=1}^{i-1} 2s_l^z \right) s_i^+ \quad (\text{A.14a})$$

$$n_i = \frac{1}{2} - s_i^z. \quad (\text{A.14b})$$

Appendix B

Dimension of spin multiplets

In this appendix we prove Eq. (2.57), as stated in Section 2.3.3:

$$D_{L,S,M} = \frac{2(2S+1)}{L+2S+2} \binom{L}{\frac{L}{2}+S}, \quad (\text{B.1})$$

where $D_{L,S,M}$ is the dimension of the fixed- (S, M) subspace of the Hilbert space of L spins- $\frac{1}{2}$ ($D_{L,S,M}$ does not depend on M).

Recall that irreducible representations of the $\mathfrak{su}(2)$ algebra are classified according to the eigenvalue $s(s+1)$ taken by the quadratic Casimir operator \mathfrak{s}^2 on them, where the “spin” $s \in \{0, \frac{1}{2}, 1, \frac{3}{2}, \dots\}$ is a non-negative half-integer. The unique representation corresponding to a given value s is $(2s+1)$ -dimensional, and we denote it by $\Gamma(s)$.

Representations form a product algebra with the outer operations

$$\begin{aligned} \Gamma(s_1) \oplus \Gamma(s_2) &: (A, B) \mapsto \Gamma(s_1)(A) \oplus \Gamma(s_2)(B), \\ \Gamma(s_1) \otimes \Gamma(s_2) &: (A, B) \mapsto \Gamma(s_1)(A) \otimes \Gamma(s_2)(B), \end{aligned}$$

where in both cases A and B are $\mathfrak{su}(2)$ elements, $\Gamma(s_k)(\cdot)$ their representatives matrices and \oplus, \otimes on the right-hand sides the direct sum and product of matrices, explicitly given by

$$A \oplus B = \begin{pmatrix} A & \mathbb{O} \\ \mathbb{O} & B \end{pmatrix}$$

and

$$\begin{pmatrix} a_{11} & \cdots & a_{1r} \\ \vdots & \ddots & \vdots \\ a_{r1} & \cdots & a_{rr} \end{pmatrix} \otimes B = \begin{pmatrix} a_{11}B & \cdots & a_{1r}B \\ \vdots & \ddots & \vdots \\ a_{r1}B & \cdots & a_{rr}B \end{pmatrix}.$$


It is well-known (see *e.g.* [70]) that the combination of two irreducible representations of $\mathfrak{su}(2)$ splits into irreducible direct summands according to the Clebsch–Gordan series

$$\Gamma(s_1) \otimes \Gamma(s_2) = \Gamma(|s_1 - s_2|) \oplus \Gamma(|s_1 - s_2| + 1) \oplus \cdots \oplus \Gamma(s_1 + s_2),$$

and in particular

$$\Gamma(S) \otimes \Gamma(1/2) = \begin{cases} \Gamma(S - 1/2) \oplus \Gamma(S + 1/2) & \text{if } S \neq 0 \\ \Gamma(1/2) & \text{otherwise.} \end{cases}$$

In other words, whenever we add a new spin- $\frac{1}{2}$ particle to an S -multiplet, two new ones are generated in its place, unless $S = 0$. This can be visualized as a branching process with a boundary:



The labels in the circles count how many distinct (L, S) multiplets are created by combining L spins- $\frac{1}{2}$ into a total spin S , or in other words, the degeneracy of the fixed- (L, S, M) subspace (for each $M \in \{-S, -S + 1, \dots, S\}$). For example, from the above figure we can immediately tell that

$$\Gamma(1/2)^{\otimes 4} = 2\Gamma(0) + 3\Gamma(1) + \Gamma(2), \quad (\text{B.2})$$

with the integer coefficients of $\Gamma(S)$ quantifying how many orthogonal spin- S states with a given magnetization M one can construct out of four spins- $\frac{1}{2}$.

The values of the labels correspond to the number of possible paths from the “origin” $(L, S) = (1, \frac{1}{2})$ to the elected (L, S) site in the “boundary-restricted Pascal triangle” [147] depicted above. We now show that this number is given by Eq. (B.1) using a version of the standard method of images.

Proof. Consider the origin $(L, S) = (1, \frac{1}{2})$ as a “source of trajectories” and imagine that there were no boundary. We denote the weight carried by these unrestricted trajectories by $P_{L,S}$ and assign unit weight to the origin, $P_{1,1/2} = 1$. The splitting rule

$$P_{L,S} = P_{L-1,S-1/2} + P_{L-1,S+1/2} \quad (L \geq 2) \quad (\text{B.3})$$

(with $P_{1,S}$ understood to be zero for all $S \neq \frac{1}{2}$) would then result in the ordinary Pascal triangle,

$$P_{L,S} = \binom{L-1}{\frac{L}{2}-S}. \quad (\text{B.4})$$

Now add a source of “anti-trajectories” at site $(1, -\frac{3}{2})$, *i.e.* a node with a weight $\bar{P}_{1,-3/2} = -1$ and splitting with the same rules as $P_{L,S}$. This obviously generates an “anti-Pascal triangle”,

$$\bar{P}_{L,S} = -\binom{L-1}{\frac{L}{2}-S-2}. \quad (\text{B.5})$$

With these prescriptions, we see that a trajectory and an anti-trajectory with opposite weight meet at $(L, S) = (3, -\frac{1}{2})$, giving a zero total weight on that site. The same holds for all sites lying on the midway line $S = -\frac{1}{2}$ between the source and the anti-source. This means that we have created a boundary at $S = -\frac{1}{2}$, ensuring that the sites on the $S = 0$ line receive no contributions from nodes to their left. But then, summing the weights of the trajectories and those of the anti-trajectories gives us the correct weight of the boundary-restricted Pascal triangle at any site (L, S) , $S \geq 0$. We therefore obtain

$$D_{L,S,M} = P_{L,S} + \bar{P}_{L,S} = \binom{L-1}{\frac{L}{2}-S} - \binom{L-1}{\frac{L}{2}-S-2}. \quad (\text{B.6})$$

In order to complete the proof, we recall two trivial properties of the binomial coefficient:

$$1) \quad \binom{n}{k} = \frac{n}{k} \binom{n-1}{k-1}; \quad (\text{B.7})$$

$$2) \quad \binom{n}{k} = \frac{n+1-k}{k} \binom{n}{k-1}. \quad (\text{B.8})$$

First, we conveniently rewrite the first term of Eq. (B.6) as

$$\binom{L-1}{\frac{L}{2}-S} = \binom{L-1}{\frac{L}{2}+S-1}, \quad (\text{B.9})$$

Then, using property (2) twice on the second term of the same equation, we have

$$\binom{L-1}{\frac{L}{2}-S-2} = \binom{L-1}{\frac{L}{2}+S+1} = \frac{(L-2S-2)(L-2S)}{(L+2S+2)(L+2S)} \binom{L-1}{\frac{L}{2}+S-1}. \quad (\text{B.10})$$

Finally, subtracting Eq. (B.10) from Eq. (B.9) we get

$$D_{L,S,M} = \frac{4(2S+1)L}{(L+2S+2)(L+2S)} \binom{L-1}{\frac{L}{2}+S-1} = \frac{2(2S+1)}{L+2S+2} \frac{L}{\frac{L}{2}+S} \binom{L-1}{\frac{L}{2}+S-1}, \quad (\text{B.11})$$

and by applying property (1) to the right-hand side we recognize Eq. (B.1). \square

As a final remark, one can show that the following asymptotic behavior holds for large L and fixed α :

$$\binom{L}{\alpha L+s} = \frac{e^{LH(\alpha)-\beta \log \frac{\alpha}{1-\alpha}}}{\sqrt{2\pi\alpha(1-\alpha)L}} \left[1 + O\left(\frac{1}{L}\right) \right], \quad (\text{B.12})$$

with $H(\alpha) = -\alpha \log \alpha - (1-\alpha) \log(1-\alpha)$. Applied to our formula ($\alpha = 1/2$), this gives us the asymptotic estimate

$$D_{L,S,M} = \frac{4(2S+1)}{\sqrt{2\pi}} \frac{2^L}{L^{3/2}} \left[1 + O\left(\frac{1}{L}\right) \right], \quad (\text{B.13})$$

to be compared with the fixed- M sector dimension

$$D_{L,M} = \binom{L}{\frac{L}{2}+M} = \sqrt{\frac{2}{\pi}} \frac{2^L}{L^{1/2}} \left[1 + O\left(\frac{1}{L}\right) \right], \quad (\text{B.14})$$

with an asymptotic gain

$$\frac{D_{L,S,M}}{D_{L,M}} \sim \frac{2(2S+1)}{L} \quad (\text{B.15})$$

at large L . For a numerical comparison, see Table 2.2 in the main text.

As a word of caution, notice that only knowing the Hilbert space dimension is not enough to draw conclusions about the effective memory requirement for *e.g.* storing¹ an operator acting on that space, because one has to account for the *sparsity* (fraction of zero-valued matrix elements) of the corresponding representative matrix as well. Local operators are (by definition) very sparse on the canonical basis, while they are not guaranteed to be sparse on a generic basis. In the case of a tree basis, the operators do remain sparse, but with a higher density of nonzero elements.

For the random-bond Heisenberg model studied in Section 2.4, the memory requirement stays favorable for the tree basis on all considered system sizes $L \leq 26$, but the

¹The true memory bottleneck actually does not come from storing the matrix in RAM, but rather from the LU factorization (or equivalent) required for its diagonalization, which is required when searching for mid-spectrum eigenstates as the less memory-heavy iterative solvers (such as the Lanczos algorithms) fail to converge in regions of the spectrum with extremely small level spacing, which make the problem ill-conditioned (see also Ref. [206]). While a precise estimate of the memory complexity of the algorithm is very hard to give, one expects it to depend on the number of nonzero elements of the matrix in question, so a sparsity analysis remains appropriate.




Figure B.1: Comparison between the representative matrices of a random-bond Heisenberg model (Eq. (2.30)) on L particles when restricted to the $(S, M) = (0, 0)$ sector of the Hilbert space via a tree basis construction (blue circles) and when instantiated in the canonical basis with only the $M = 0$ states retained (orange squares). (Left) The total number of matrix elements, obtained by squaring the Hilbert space dimensions $D_{L,S,M}$ and $D_{L,M}$ respectively. The lines are the squares of the asymptotic results (B.13) and (B.14). (Middle) The fraction of nonzero matrix elements of the models. The lines represent exponential fits $f(L) = \alpha e^{-\beta L}$, with fitted exponents $\beta_{\text{tree}} = 0.447(5)$ and $\beta_{\text{canon}} = 0.596(2)$. (Right) The actual number of nonzero matrix elements. The lines are obtained by combining the ones in the left and middle figures.

memory gain for storing the Hamiltonian in the $(S, M) = (0, 0)$ sector with a tree basis is only about 50% at $L = 20$, compared to storing a canonical-basis matrix restricted to the $M = 0$ sector (*i.e.* without spin restriction), see Fig. B.1. Asymptotic analysis suggests that the canonical M -restricted basis actually becomes preferable around $L \approx 28$. However, this analysis assumes that one is able to directly instantiate the spin operators in the M -restricted basis, which is not standardly done in the canonical case.²

In conclusion, for all practical purposes, the tree basis construction appears to always outperform (memory-wise) the canonical one in the range of interest.

²The simplest way to M -restrict a canonical matrix is to first instantiate the full operator, which is done in terms of Kronecker products, and then “masking out” the rows and columns corresponding to states with the wrong magnetization. Building the full operator requires a prohibitive amount of memory compared to the restricted one.

Statistical properties of fusion trees

In this appendix, we study the properties of *random* fusion trees. Unlike the actual SDRG trees introduced in Section 2.4.2, these are obtained by fusing neighboring spins together at random, ignoring the values of the (bare or renormalized) couplings between them. In other words, only the information about their spatial arrangement is retained, while the coupling structure is “disorder-averaged” away. This simplification is instrumental in allowing us to reach some analytical results.

C.1 Distribution of nearest-neighbor graph distances

We want to prove the claim (2.79) in the main text,

$$P(l) = \frac{3}{4} \left(\frac{2}{3}\right)^l, \quad \text{for } L \rightarrow \infty, \tag{C.1}$$

where $P(l)$ is the distribution of the random variable $l_{i,i+1}$, namely the graph distance of two neighboring spins in a generic SDRG tree.

To this end, we consider the ensemble of trees constructed by taking a chain of L spins and fusing them all together, two neighbors at a time. After each fusion, the chain effectively shrinks by one site, and the neighbor structure gets updated accordingly. This is an approximation of the SDRG procedure where we completely neglect the detailed structure of the J couplings.

More precisely, let a tree be described by a sequence of fusions (i_1, \dots, i_{L-1}) , where i_k means that we are fusing, at the k -th step, the pair $(i_k, i_k + 1)$ (with periodic boundary conditions). In order to emulate the SDRG algorithm, we sample the sequence of fusions uniformly randomly among the $L!$ possible $(L - 1)$ -permutations of $(1, \dots, L)$. This results in a *biased* distribution on the set of all binary trees, with “taller” trees being less likely.

Now take a generic pair $(i, i + 1)$ in a given tree, and suppose that their common block-spin descendant was created at the $(k + 1)$ -th step of the tree construction. All the fusions taking place after that step are irrelevant for determining $l_{i,i+1}$, whereas each of the k previous ones may contribute either 0 or 1 to such distance. In fact, the distance contributed by the j -th fusion is a Bernoulli random variable with success probability $p_j = \frac{2}{L-j}$, because the distance between i and $(i + 1)$ only increases if either one of their descendants is picked out of the $L-j$ possible spins at that step. Moreover, the contributions are uncorrelated since all free indices are sampled with equal probability regardless of the previous history of the tree construction.

Therefore we have

$$l_{i,i+1} = l^{(k)} = 2 + x_1 + x_2 + \dots + x_k \quad (\text{C.2})$$

where

$$x_j = \begin{cases} 1 & \text{with probability } p_j = \frac{2}{L-j} \\ 0 & \text{" " } 1 - p_j \end{cases} \quad (\text{C.3})$$

and the (k) superscript serves as a reminder that our random variable is now being conditioned on k .

Let us compute the cumulant generating function for $l^{(k)} - 2$:

$$\langle e^{-s(l^{(k)}-2)} \rangle = \prod_{j=1}^k (1 - p_j + p_j e^{-s}), \quad (\text{C.4})$$

the logarithm of which is

$$\log \langle e^{-s(l^{(k)}-2)} \rangle = \sum_{j=1}^k \log \left(1 + \frac{1}{L} \frac{2}{1-j/L} (e^{-s} - 1) \right). \quad (\text{C.5})$$

By defining $x = j/L$, $\alpha = k/L$, and taking $L \rightarrow \infty$, we have

$$\begin{aligned} \log \langle e^{-s(l^{(k)}-2)} \rangle &\sim \int_0^\alpha dx \frac{2}{1-x} (e^{-s} - 1) \\ &= 2(1 - e^{-s}) \log(1 - \alpha) \end{aligned} \quad (\text{C.6})$$

up to $O(1/L)$ terms.

Now notice that $(k + 1)$, in our ensemble, is uniformly distributed between 1 and $L - 1$, as it corresponds to the position of index i in the tuple (i_1, \dots, i_{L-1}) . We can then get rid of the k -conditioning by averaging over $\alpha \in [0, 1]$. This gives

$$\begin{aligned} \langle e^{-s(l-2)} \rangle &= \int_0^1 d\alpha (1 - \alpha)^{2(1-e^{-s})} \\ &= \frac{1}{3 - 2e^{-s}}, \end{aligned} \quad (\text{C.7})$$

and by expanding the denominator in a geometric series, we get

$$\langle e^{-sI} \rangle = \frac{1}{3} \left[e^{-2s} + \frac{2}{3} e^{-3s} + \left(\frac{2}{3} \right)^2 e^{-4s} + \dots \right]. \quad (\text{C.8})$$

Inverting the Laplace transform results in Eq. (2.79).

C.2 Size of the block spins

We now set out to determine the average size of the support of a randomly chosen block-spin operator for a random SDRG tree state. This amounts to estimating the number of leaves which connect to a node picked uniformly randomly from the set of non-leaf nodes in a generic fusion tree.

To this end, it is convenient to introduce an alternative (but equivalent) construction for our random ensemble. Consider a single node, and start by attaching two children nodes to it, one to the left and one to the right. We can see this as a “splitting” step for the original node. Now pick with equal probability either one of the resulting leaves and perform the same kind of splitting. Iterate the procedure for a total of $(L - 1)$ times, such that the final number of leaves is L . The leaves are spacially ordered by the order relation induced in an obvious way by the distinction of left- and right-children. In this way we obtain a binary tree whose geometry is compatible with an SDRG tree. We can call this the “fission tree” ensemble.

We are now going to prove by induction that the fission tree and fusion tree ensembles are equivalent ¹.

Suppose that the above claim holds after the $(k - 1)$ -th splitting, that is to say, for the ensembles of k -leaved fission and fusion trees. Now, when constructing a *fusion* tree on $(k + 1)$ leaves, after the first fusion we are left with an effective k -leaved tree. In order to prove the claim it is then enough to show that the first fusion does not spoil the ensemble equivalence. By definition of the fusion tree ensemble, it is the case that each one of the initial $(k + 1)$ leaf pairs has the same probability of being fused at the first step, which means that every one of the k effective leaves after the first step has the same likelihood of being the one resulting from the fusion. Therefore, upon reversing the “time direction” we see that if we allow all the k leaves to split with the same probability, both fission and fusion trees on $(k + 1)$ leaves are sampled with the same distribution, and the inductive step is completed. It also holds trivially that the two ensembles coincide when $k = 1$, providing the basis of the induction.

In light of this, it is possible to assign to each node of a tree the step at which it was split. For instance, the root will always be labeled by 1, and the maximum label will be

¹I am grateful to user Misha Lavrov for suggesting this proof as part of a reply to a question on the math.stackexchange.com website.

$L - 1$ (note that, similarly to the case of the “fusion labeling” in Appendix C.1, this labeling is not uniquely defined). Now fix $k \in \{1, \dots, L - 1\}$ and consider the node labeled by k . Introduce the variable t to measure the number of fissions occurring after the k -th one, $t \in \{0, \dots, L - k - 1\}$, and call $N(t)$ the total number of leaves which affect the state of the initial node at “time” t . Since every fission can only increment N by 1 at time $(t + 1)$ if one of the $N(t)$ leaves is picked for the fission, we have the stochastic recursion equation

$$N^{(k)}(t + 1) = N^{(k)}(t) + B[p^{(k)}(t)], \quad (\text{C.9})$$

where $B[p]$ is a Bernoulli variable with success probability p , and $p^{(k)}(t) = \frac{N(t)}{k+t+1}$. The initial condition must be set to $N^{(k)}(0) = 2$.

This equation is hard to treat due to the $N(t)$ -dependence hidden inside $p^{(k)}(t)$, but it is linear, and therefore easily solved in the expectation values:

$$\overline{N^{(k)}(t + 1)} = \overline{N^{(k)}(t)} \left(1 + \frac{1}{k + t + 1} \right), \quad (\text{C.10})$$

where we used $\overline{B[p]} = p$. By iterating and simplifying the product on the right hand side, and then looking at the final time, we get

$$\overline{N^{(k)}} = \overline{N^{(k)}(L - k - 1)} = \frac{2L}{k + 1}. \quad (\text{C.11})$$

This is the average number of ancestors of a node that was split at the k -th fission step. In order to answer our initial question — what is the average number of ancestor elementary spins of a random non-leaf node —, we simply take the average on all possible values of k . This yields

$$\overline{N} = \frac{1}{L} \sum_{k=1}^{L-1} \overline{N^{(k)}} = 2 (\log L + \gamma - 1) + O\left(\frac{1}{L}\right) \quad (\text{C.12})$$

(with $\gamma \approx 0.5772$ the Euler–Mascheroni constant), showing that the block spins have on average unbounded support in space.

Correlated cost functions

Both the impurity band model and the random energy model studied in Chapter 4 have the undesirable properties of being very unrealistic ensembles when it comes to optimization problems. In real-world applications, including the “optimization” of physical systems (namely, finding low-lying states of a Hamiltonian), cost functions are typically characterized by some *locality* structure, in the sense that the cost function can be written as a sum of terms that involve a bounded number of degrees of freedom. In this appendix, we take a quick look at this kind of more realistic cost functions or “potentials” in the context of the PT algorithm.

More specifically, suppose that a problem (physically, a system) is comprised of n elementary and discrete constituents. These can be n variables of a constraint satisfaction problem, n particles, *etc.* Without loss of generality, we can suppose they are binary variables (spins- $\frac{1}{2}$) $z = (z_0, \dots, z_{n-1})$, with the computer scientist’s notation $z_i \in \{0, 1\}$. We call a bit-valued function of such variables a *clause*, $C : \{0, 1\}^k \rightarrow \{0, 1\}$. The number k of variables nontrivially involved in the clause is called its *arity* or *locality*. We say that a clause is *satisfied* by a bit string σ if $C(\sigma) = 0$, and *violated* otherwise.

We say that a cost function (Hamiltonian) H_p is *k-local* if it can be written in the form

$$H_p = \sum_{\lambda=0}^{m-1} w_\lambda C_\lambda \tag{D.1}$$

with the arities of all clauses $\{C_\lambda\}_\lambda$ upper bounded by k . Here w_λ are constant real weights, which we can take to be positive (if $w_\lambda < 0$ we can simply negate $C_\lambda \mapsto 1 - C_\lambda$ and disregard the additive constant w_λ). We can also assume that the ground level of H_p is zero.

The interaction structure of Eq. (D.1) can be represented as a “factor graph”, *i.e.* a bipartite simple graph with one node per each clause and variable ($n + m$ total), and edges




Figure D.1: A graphical representation of a Hamiltonian of the form Eq. (D.1), with $n = 5$, $m = 4$. Here, clauses C_1 , C_2 and C_3 are binary, while C_0 is ternary, so we say that H is 3-local.

connecting variables with all clauses in which they participate (Fig. D.1). For a k -local Hamiltonian, each clause node has degree less than or equal to k .

When a cost function is sufficiently local, say $k \in o(n)$, the energy levels that it generates are correlated, *i.e.* the two-level probability $P(H_p(z), H_p(z'))$ fails to factorize in general. Intuitively, this is because the locality of the clauses implies that single bit-flips cannot change the total cost by too much, so one expects that nearby (in Hamming space) configurations will have similar costs. This feature is totally absent from the (highly nonlocal) QIBM and QREM — indeed, the lack of correlation is precisely why they were defined in the first place. There is a rather illuminating way of introducing correlations in the QREM, by regarding it as a limiting case of a k -local Hamiltonian with $k \in o(n) \cap \omega(n^0)$.

Namely, consider the p -spin model (the name p rather than k is standard in this case)

$$H^{(p)} = \sum_I J_I \sigma_{i_1} \dots \sigma_{i_p} \quad (\text{D.2})$$

with the index $I = \{i_1, \dots, i_p\}$ running over all possible p -subsets of $\{0, \dots, n-1\}$ (*i.e.* we require $i_1 < \dots < i_p$), and by convention we use “spin” variables $\sigma_i = 1 - 2z_i \in \{\pm 1\}$. The couplings J_I are chosen randomly from a zero-centered Gaussian distribution with variance $p!/2n^{p-1}$; since $H^{(p)}$ is the sum of $\Theta(n^p)$ terms, this ensures that the spectrum is asymptotically almost surely contained in $[-E_0, E_0]$ with $E_0 \in \Theta(n)$ (extensivity of the energy bandwidth). Then, if $p \rightarrow \infty$ while $p \ll n$, it is possible to show [81] that the energy levels of $H^{(p)}$ become uncorrelated, *i.e.* the Hamiltonian tends to the random

energy model.¹

Consider however the Hamiltonian (D.2) for a *finite* value of p . In this case, the model contains nontrivial correlations between the energy levels, as can be seen by computing their covariance.² It is not hard to calculate this quantity for the chosen Gaussian ensemble (the (p) superscript is omitted for convenience):

$$\begin{aligned}
\mathbb{E}[H(\sigma)H(\tilde{\sigma})] &= \sum_{I'} \mathbb{E}[J_I J_{I'}] \sigma_{i_1} \cdots \sigma_{i_p} \tilde{\sigma}_{i'_1} \cdots \tilde{\sigma}_{i'_p} \\
&= \frac{p!}{2n^{p-1}} \sum_I \sigma_{i_1} \cdots \sigma_{i_p} \tilde{\sigma}_{i_1} \cdots \tilde{\sigma}_{i_p} \\
&= \frac{n}{2} \left(\frac{1}{n} \sum_i \sigma_i \tilde{\sigma}_i \right)^p \\
&= \frac{n}{2} \left(1 - \frac{2}{n} |\sigma - \tilde{\sigma}|_{\text{H}} \right)^p \\
&\sim \frac{n}{2} e^{-x_{\sigma\tilde{\sigma}}/\xi} \quad (x_{\sigma\tilde{\sigma}} \ll 1).
\end{aligned} \tag{D.4}$$

In the first passage we used the fact that $\mathbb{E}[J_I J_{I'}] = \frac{p!}{2n^{p-1}} \delta_{II'}$, and in the second one we recognized that the sum over I , once transformed into a sum over all p -ples (i_1, \dots, i_p) (this can be done by absorbing the $p!$ prefactor), actually factorizes into p identical sums; then, we recast the overlap $q_{\sigma\tilde{\sigma}} = \sum_i \sigma_i \tilde{\sigma}_i$ between the bit strings in terms of their Hamming distance, $q_{\sigma\tilde{\sigma}} = n - |\sigma - \tilde{\sigma}|_{\text{H}}$, and finally we defined the fractional Hamming distance $x_{\sigma\tilde{\sigma}} = \frac{1}{n} |\sigma - \tilde{\sigma}|_{\text{H}}$ and took the small- $x_{\sigma\tilde{\sigma}}$ limit, which allows us to define the *correlation length* $\xi = 2/p$. Notice how ξ decreases with p and correctly vanishes in the REM ($p = \infty$).

Similar considerations can be extended to some problem classes of interest in optimization theory. In particular, we consider a class of problems defined by a local cost function of the form

$$H(\sigma) = \sum_{a=0}^{m-1} C_a(\sigma) \tag{D.5}$$

defined on the Boolean n -hypercube. This is a simplified version of Eq. (D.1) with all weights set to one. Typical examples include *constraint satisfaction problems* (CSPs) —

¹The reason for the $p!/2$ factor in the variance of J_I is precisely to make the variance of the marginal energy distribution $p(E)$ asymptotic to $n/2$, which is the standard choice for the REM.

²Recall that the *linear correlation coefficient* between variables X and Y is defined as the ratio between their covariance and the product of their standard deviations:

$$r(X, Y) = \frac{\mathbb{E}[XY] - \mathbb{E}[X]\mathbb{E}[Y]}{\sqrt{\text{Var}[X]\text{Var}[Y]}}. \tag{D.3}$$

where each clause represents a single constraint on a set of variables — in which all constraints have the same importance.

Unlike computer scientists, who tend to focus on worst-case scenario instances for the definition of the complexity of a problem class such as (D.5), statistical physicists are more interested in average or typical properties of *problem ensembles*, defined by specifying a probability distribution over the problem class. In analogy with the p -spin model, we define an ensemble by picking every clause C_a independently with the same distribution $P(C)$. This makes the evaluation $C_a(\sigma)$ a random variable for each a with respect to the disorder distribution, for any *fixed* σ (i.e. σ is not itself a random variable — only the evaluated clause).

For simplicity, we ask that $P(C)$ be *unbiased* and *isotropic*. Unbiasedness means that the distribution should not favor any bit string over any other:

$$\mathbb{P}[C(\sigma) = 1] = p_0 \quad \forall \sigma \in \{0, 1\}^n. \quad (\text{D.6})$$

Then obviously

$$\mathbb{E}[H(\sigma)] = mp_0 \quad (\text{D.7})$$

for any bit string σ . The second requirement, isotropy, is a statement about the two-point function. Fix a pair of bit strings $(\sigma, \tilde{\sigma})$, and call

$$\pi(\sigma, \tilde{\sigma}) = \mathbb{P}[C(\sigma) = C(\tilde{\sigma}) = 1] \quad (\text{D.8})$$

the probability that a random clause is simultaneously violated by both of them. We demand that this probability only depend on the Hamming distance between the two strings:

$$\pi(\sigma, \tilde{\sigma}) = p_0 p_1(x_{\sigma\tilde{\sigma}}), \quad (\text{D.9})$$

where the fractional Hamming distance $x_{\sigma\tilde{\sigma}}$ was defined previously. We included a p_0 prefactor for convenience: in this way, by elementary probability theory, $p_1(x)$ represents the probability that a clause is violated by some fixed bit string $\tilde{\sigma}$, *conditioned* on the fact that it is already violated by some other fixed bit string σ , where $x_{\sigma\tilde{\sigma}} = x$.

The problem ensemble we have defined includes models such as the random versions of MAX- k -SAT and MAX- k -XORSAT, which are characterized by cost functions whose clauses involve uniformly random k -subsets of the literal space. In the former, a clause is violated if all the involved variables have value 0; in the other, if an even number of them have value 0. The underlying problem classes are both NP-hard.

We are now ready to derive a general formula for the correlation. We first compute the

two-point function

$$\begin{aligned}
\mathbb{E}[H(\sigma)H(\tilde{\sigma})] &= \sum_{a=0}^{m-1} \sum_{b=0}^{m-1} \mathbb{E}[C_a(\sigma)C_b(\tilde{\sigma})] \\
&= \sum_{a=0}^{m-1} \sum_{b=0}^{m-1} \mathbb{P}[C_a(\sigma) = C_b(\tilde{\sigma}) = 1] \\
&= \sum_{a=0}^{m-1} \sum_{b=0}^{m-1} [\delta_{ab}p_0p_1(x_{\sigma\tilde{\sigma}}) + (1 - \delta_{ab})p_0^2] \\
&= mp_0p_1(x_{\sigma\tilde{\sigma}}) + m(m-1)p_0^2,
\end{aligned} \tag{D.10}$$

(where we used clause independence and the definition of $p_1(x)$ to write the third line and subtract the disconnected part $\mathbb{E}[H(\sigma)]\mathbb{E}[H(\tilde{\sigma})] = m^2p_0^2$ (see Eq. (D.7)) to obtain the covariance:

$$\text{Cov}[H(\sigma), H(\tilde{\sigma})] = mp_0(p_1(x_{\sigma\tilde{\sigma}}) - p_0). \tag{D.11}$$

Finally, we normalize by the variance (obtained by setting $x_{\sigma\tilde{\sigma}} = 0$ in Eq. (D.11)) to obtain the correlation coefficient:

$$r(\sigma, \tilde{\sigma}) = \frac{\text{Cov}[H(\sigma), H(\tilde{\sigma})]}{\text{Var}[H(\sigma)]} = \frac{p_1(x_{\sigma\tilde{\sigma}}) - p_0}{p_1(0) - p_0}. \tag{D.12}$$

We thus see that for our problem class Eq. (D.5), the correlation coefficient (and hence the correlation length) is fully determined by the knowledge of the quantity p_0 and the function $p_1(x)$, defined respectively in Eqs. (D.6) and (D.9).

Examples

- a) In random MAX- k -SAT, the probability of violating a random clause is given by the probability that none of the k literals agrees with the assignment, which is

$$p_0^{\text{SAT}} = \frac{1}{2^k}. \tag{D.13}$$

When conditioning on $x = x_{\sigma\tilde{\sigma}}$, the probability to violate C with $\tilde{\sigma}$, given that σ violates it, amounts to the probability that all literals in C have been assigned equal value in σ and $\tilde{\sigma}$ (otherwise, at least one variable in $\tilde{\sigma}$ would agree with at the corresponding clause literal, satisfying the clause). From a dual perspective, the probability that all k literals in clause C are picked from the $(1-x)n$ variables which are common to σ and $\tilde{\sigma}$ is

$$p_1^{\text{SAT}}(x) = \frac{\binom{(1-x)n}{k}}{\binom{n}{k}} \sim (1-x)^k \quad (n \gg 1). \tag{D.14}$$




Figure D.2: Linear correlation coefficient $r(x) = \text{Cov}[H(\sigma)H(\tilde{\sigma})] / \text{Var}[H(\sigma)]$ between energy levels in random MAX- k -SAT (blue lines) and random MAX- k -XORSAT (red lines) as a function of the fractional Hamming distance x between σ and $\tilde{\sigma}$ (the only relevant variable due to the assumptions of clause unbiasedness and isotropy). The results are plotted for $k = 3$ (full lines) and $k = 4$ (dashed lines).

Therefore, using Eq. (D.12), we find that the energy density correlation is

$$r^{\text{SAT}}(x) = \frac{(1-x)^k - 2^{-k}}{1 - 2^{-k}}. \quad (\text{D.15})$$

The correlation length is $\xi^{\text{SAT}} = 1/k$.

- b) In random MAX- k -XORSAT, where clauses are satisfied or violated based on their parity, the probability of violating a random clause is always

$$p_0^{\text{XORSAT}} = \frac{1}{2} \quad (\text{D.16})$$

regardless of k . Upon conditioning, a similar reasoning as the one above leads to the conclusion that in order for $\tilde{\sigma}$ to violate a σ -violated clause, the number of clause literals which do not agree between the two assignments must be an even number. The corresponding probability can be written as

$$\begin{aligned} p_1^{\text{XORSAT}}(x) &= \binom{n}{xn}^{-1} \sum_{s \text{ even}} \binom{k}{s} \binom{n-k}{xn-s} \\ &\sim \frac{1}{2} \left[1 + (1-2x)^k \right], \end{aligned} \quad (\text{D.17})$$

where we used the identity

$$\sum_{s \text{ even}} \binom{k}{s} q^s = \frac{1}{2} [(1 - q)^k + (1 + q)^k]. \quad (\text{D.18})$$

As a result, from Eq. (D.12) the correlation simply reads

$$r^{\text{XORSAT}}(x) = (1 - 2x)^k. \quad (\text{D.19})$$

The correlation length is $\xi^{\text{XORSAT}} = 2/k$.

The correlation coefficients for the two ensembles are displayed in Fig. D.2.

We conclude by pointing out that the existence of local correlations in realistic problems such as CSPs is precisely what accounts for the possibility of creating successful classical algorithms to solve them. In general, problems without a recognizable structure cannot be tackled by anything short of random or brute-force search, because there is no heuristic or insight to guide our efforts. On the contrary, when problems admit a “good” representation *e.g.* in terms of factor graphs, several heuristics are available, often resulting in much more efficient algorithms [117]. As a consequence, in a realistic setting we expect comparisons between population transfer and “classical optimality” to be more difficult to draw. We refer to Section 4.4 for further discussion.

Uniformity in Hamming space

In Section 4.3.3, one of the proposed ways to assess the effectiveness of PT is to understand to what extent sampling of the resonant space is performed “fairly”, *i.e.* without bitstring bias. This corresponds to asking how uniformly the wavefunction is spreading over the microcanonical shell Ω as a consequence of time evolution. Here, “uniformly” means that the probability distribution ought to explore the entire space Ω , and not remain concentrated in some corner—*e.g.* the vicinity of the initial state.

It is well-known that the Shannon entropy $S[p]$ is maximized, with respect to all pdfs defined on a given domain, precisely when p is the uniform distribution on that domain, so one may think of using S as a proxy for well-spreadness. However, S is in all respects insensitive to the *spatial* structure of p : roughly speaking, it only counts on how many states p is supported, but without care for their mutual Hamming distance.

What we would like to measure instead, is how efficiently PT populates states which are far apart in Hamming distance. To this end, we introduce the “repulsive potential”

$$U_{\Omega}[\psi] = \frac{\sum_{i,j \in \Omega} |\psi_i|^2 |\psi_j|^2 u_{ij}}{W_{\Omega}[\psi]^2}, \quad (\text{E.1})$$

where $W_{\Omega}[\psi] = \sum_{i \in \Omega} |\psi_i|^2$ is the total probability of sampling a bitstring inside Ω and $u_{ij} = u(|i - j|_{\text{H}})$ is a symmetric, positive-definite “two-body term” which decreases with the Hamming distance $|i - j|_{\text{H}}$ between i and j . We take it to be

$$u_{ij} = e^{-\frac{br_{ij}}{n}}, \quad (\text{E.2})$$

where $r_{ij} = |i - j|_{\text{H}}$ and b is an arbitrary positive parameter which we fix to 1.

From a physical perspective, U_{Ω} is akin to an electrostatic potential for the “charge distribution” $p_i = W_{\Omega}^{-1} |\psi_i|^2$ defined on Ω . The denominator ensures that this distribution is properly normalized, so that U_{Ω} does not depend on the behavior of $|\psi|^2$ outside the

resonant space in the sense that a uniform spillage $|\psi_i|^2 \mapsto \lambda|\psi_i|^2, \forall i \in \Omega$ (with arbitrary λ) does not change the value of U_Ω .

It is easy to see that U_Ω is maximized when all the “charge” is clumped together, $p_i = \delta_{i,z_0}$, which corresponds to the fully localized case, $\Gamma = 0$. Then U_Ω achieves its maximum, $U_\Omega^{\text{loc}} = 1$. Indeed, trying to spread the probability to another site, $p_i = (1 - \varepsilon)\delta_{i,z_0} + \varepsilon\delta_{i,z_1}$, results in a lower potential $U_\Omega = 1 - 2\varepsilon(1 - \varepsilon)(1 - u_{z_0,z_1}) < 1$.

On the opposite end, when p_i is uniform on Ω we expect the functional to be minimal. Indeed, call U_0 its value,

$$U_0 = \frac{1}{|\Omega|^2} \sum_{i,j \in \Omega} u_{ij}, \quad (\text{E.3})$$

and consider an arbitrary perturbation of the uniform distribution, $p_i = p_0 + \delta p_i$, where $p_0 = 1/|\Omega|$ and the perturbation is *nonspilling*,

$$\delta W_\Omega = \sum_{i \in \Omega} \delta p_i = 0. \quad (\text{E.4})$$

This is generic as nonzero values of δW_Ω can be decomposed into uniform spillage, which does not affect U_Ω , composed with a nonspilling perturbation, namely $\delta p_i = \frac{\delta W_\Omega}{|\Omega|} + \delta p'_i$ where $\delta p'_i$ is nonspilling.

The functional is now written in the form

$$U_\Omega[\psi] = U_0 + \frac{2}{|\Omega|} \sum_{i,j \in \Omega} \delta p_i u_{ij} + \sum_{i,j \in \Omega} \delta p_i \delta p_j u_{ij}. \quad (\text{E.5})$$

We now observe that, *on average*, the Ω space represents the Hamming cube uniformly (in distribution), namely, we can operate the substitution

$$\sum_{i,j \in \Omega} = \frac{|\Omega|^2}{D^2} \sum_{i,j} \quad (\text{E.6})$$

when we are interested in the average properties of U_Ω (here and throughout, an unspecified summation domain always refers to the entire Hamming cube $\{0, \dots, D-1\}$). This is because the REM has totally uncorrelated energy levels, so the two-point indicator factorizes: $\mathbb{E} [\chi_\Omega(i)\chi_\Omega(j)] = \mathbb{E} [\chi_\Omega(i)]\mathbb{E} [\chi_\Omega(j)] = |\Omega|^2/D^2$.

Therefore, when we rewrite the sum in the second term at the r.h.s. of Eq. (E.5) as

$$\sum_{i,j \in \Omega} \delta p_i u_{ij} = \sum_{i \in \Omega} \delta p_i \sum_{j \in \Omega} u_{ij}, \quad (\text{E.7})$$

we recognize that the inner sum $\sum_{j \in \Omega} u_{ij} = \frac{|\Omega|}{D} \sum_j u_{ij}$ (equality in expectation values) does not actually depend on i , since $\sum_j u_{ij} = \sum_{r=0}^n \binom{n}{r} u(r)$. As a consequence, the whole term vanishes due to Eq. (E.4).

This, along with the fact that the third term at the r.h.s. of Eq. (E.5) is a quadratic form of the positive definite matrix Eq. (E.2) ¹, implies that $U_{\Omega}[\psi] \geq U_0$, with equality only in the uniform case $\delta p_i = 0$.

This indicates that U_{Ω} can be used as a measure of uniformity of the wavefunction intensity distribution. We conclude by estimating the specific value of U_0 expected of a totally extended system.

From Eqs. (E.3) and (E.6), we have

$$U_0 = \frac{1}{D^2} \sum_{i,j} e^{-\frac{br_{ij}}{n}} = \frac{1}{D} \sum_{r=0}^n e^{-\frac{br}{n}} = \left(\frac{1 + e^{-b/n}}{2} \right)^n \sim e^{-b/2}, \quad (\text{E.8})$$

in good accord with both the NEE and ergodic curves in Fig. 4.7.

¹The spectrum of $(u_{ij})_{ij}$ can be computed analytically using general properties of “exponential distance matrices” [60], showing that the minimum eigenvalue is $(1 - e^{-b/n})^n > 0$.

Bibliography

- [1] S. Aaronson. Complexity Zoo. https://complexityzoo.uwaterloo.ca/Complexity_Zoo [accessed: 2020-10-03].
- [2] D. A. Abanin, E. Altman, I. Bloch, and M. Serbyn. Colloquium: Many-body localization, thermalization, and entanglement. *Rev. Mod. Phys.*, **91**(2):021001, May 2019.
- [3] D. A. Abanin, W. De Roeck, and F. Huveneers. Exponentially slow heating in periodically driven many-body systems. *Phys. Rev. Lett.*, **115**(25):256803, Dec 2015.
- [4] D. A. Abanin and Z. Papić. Recent progress in many-body localization. *Annalen der Physik*, **529**(7):1700169, Jul 2017.
- [5] R. Abou-Chacra, D. J. Thouless, and P. W. Anderson. A selfconsistent theory of localization. *Journal of Physics C: Solid State Physics*, **6**(10):1734–1752, May 1973.
- [6] E. Abrahams, P. W. Anderson, D. C. Licciardello, and T. V. Ramakrishnan. Scaling theory of localization: Absence of quantum diffusion in two dimensions. *Phys. Rev. Lett.*, **42**(10):673–676, Mar 1979.
- [7] A. A. Abrikosov. Magnetic impurities in nonmagnetic metals. *Soviet Physics Uspekhi*, **12**(2):168–181, Feb 1969.
- [8] K. Agarwal, E. Demler, and I. Martin. $1/f^\alpha$ noise and generalized diffusion in random heisenberg spin systems. *Phys. Rev. B*, **92**(18):184203, Nov 2015.
- [9] D. Aharonov, W. van Dam, J. Kempe, Z. Landau, S. Lloyd, and O. Regev. Adiabatic quantum computation is equivalent to standard quantum computation. *SIAM Journal on Computing*, **37**(1):166–194, 2007.
- [10] M. Aizenman and S. Warzel. *Random Operators*. Graduate Studies in Mathematics. American Mathematical Society, 2015.

- [11] T. Albash and D. A. Lidar. Adiabatic quantum computation. *Rev. Mod. Phys.*, **90**(1), Jan 2018.
- [12] T. Albash, T. F. Rønnow, M. Troyer, and D. A. Lidar. Reexamining classical and quantum models for the D-Wave One processor. *Eur. Phys. J. Spec. Top.*, **224**(1):111–129, Feb 2015.
- [13] D. Z. Albert. On quantum-mechanical automata. *Phys. Lett. A*, **98**(5):249–252, 1983.
- [14] F. Alet and N. Laflorencie. Many-body localization: An introduction and selected topics. *Comptes Rendus Physique*, **19**(6):498–525, Sep–Oct 2018.
- [15] B. L. Altshuler, E. Cuevas, L. B. Ioffe, and V. E. Kravtsov. Nonergodic phases in strongly disordered random regular graphs. *Phys. Rev. Lett.*, **117**(15):156601, Oct 2016.
- [16] A. Ambainis and O. Regev. An elementary proof of the quantum adiabatic theorem. *arXiv e-prints*, Jul 2006. arXiv:quant-ph/0411152v2.
- [17] L. Amico, R. Fazio, A. Osterloh, and V. Vedral. Entanglement in many-body systems. *Rev. Mod. Phys.*, **80**(2):517, May 2008.
- [18] M. H. S. Amin. Consistency of the adiabatic theorem. *Phys. Rev. Lett.*, **102**(22):220401, Jun 2009. [Erratum: *Phys. Rev. Lett.*, **115**(3):039902, Jul 2015].
- [19] P. W. Anderson. Absence of diffusion in certain random lattices. *Phys. Rev.*, **109**(5):1492–1505, Mar 1958.
- [20] P. W. Anderson. Theory of magnetic exchange interactions: exchange in insulators and semiconductors. *Solid State Physics*, **14**:99–214, 1963.
- [21] P. W. Anderson. Spin glass I: a scaling law rescued. *Physics Today*, **41**(1):9–11, 1988.
- [22] E. Andriyash and M. H. Amin. Can quantum Monte Carlo simulate quantum annealing? *arXiv e-prints*, Mar 2017. arXiv:1703.09277 [quant-ph].
- [23] S. Arora. Polynomial time approximation schemes for euclidean traveling salesman and other geometric problems. *J. ACM*, **45**(5):753–782, Sep 1998.
- [24] S. Arora, C. Lund, R. Motwani, M. Sudan, and M. Szegedy. Proof verification and the hardness of approximation problems. *J. ACM*, **45**(3):501–555, May 1998.
- [25] C. Artiaco, F. Balducci, G. Parisi, and A. Scardicchio. Quantum jamming: Critical properties of a quantum mechanical perceptron. *arXiv e-prints*, Mar 2020. arXiv:2003.01073 [cond-mat.dis-nn].

- [26] F. Arute, K. Arya, R. Babbush, et al. Quantum supremacy using a programmable superconducting processor. *Nature*, **574**:505–510, Oct 2019.
- [27] N. W. Ashcroft and N. D. Mermin. *Solid State Physics*. Harcourt Brace College Publishers, 1st edition, 1976.
- [28] Y. Y. Atas, E. Bogomolny, O. Giraud, and G. Roux. Distribution of the ratio of consecutive level spacings in random matrix ensembles. *Physical Review Letters*, **110**(8), Feb 2013.
- [29] G. Ausiello, P. Crescenzi, G. Gambosi, V. Kann, A. Marchetti-Spaccamela, and M. Protasi. *Complexity and Approximation: Combinatorial Optimization Problems and Their Approximability Properties*. U.S. Government Printing Office, 1999.
- [30] M. Baity-Jesi, G. Biroli, and C. Cammarota. Activated aging dynamics and effective trap model description in the random energy model. *Journal of Statistical Mechanics: Theory and Experiment*, 2018(1):013301, Jan 2018.
- [31] C. L. Baldwin and C. R. Laumann. Quantum algorithm for energy matching in hard optimization problems. *Phys. Rev. B*, **97**(22):224201, Jun 2018.
- [32] C. L. Baldwin, C. R. Laumann, A. Pal, and A. Scardicchio. The many-body localized phase of the quantum random energy model. *Phys. Rev. B*, **93**(2):024202, Jan 2016.
- [33] V. Bapst, L. Foini, F. Krzakala, G. Semerjian, and F. Zamponi. The quantum adiabatic algorithm applied to random optimization problems: The quantum spin glass perspective. *Physics Reports*, **523**(3):127–205, Feb 2013.
- [34] J. H. Bardarson, F. Pollmann, and J. E. Moore. Unbounded growth of entanglement in models of many-body localization. *Phys. Rev. Lett.*, **109**(1):017202, Jul 2012.
- [35] 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, 2006.
- [36] B. Bauer and C. Nayak. Area laws in a many-body localized state and its implications for topological order. *Journal of Statistical Mechanics: Theory and Experiment*, **2013**(09):P09005, Sep 2013.
- [37] R. J. Baxter. *Exactly solved models in statistical mechanics*. American Press, London, 1982.
- [38] M. J. S. Beach, R. G. Melko, T. Grover, and T. H. Hsieh. Making trotters sprint: A variational imaginary time ansatz for quantum many-body systems. *Phys. Rev. B*, Sep 2019.

- [39] F. Becca and S. Sorella. *Quantum Monte Carlo Approaches for Correlated Systems*. Cambridge University Press, Cambridge, 2017.
- [40] A. Berman and R. J. Plemmons. *Nonnegative Matrices in the Mathematical Sciences*. Classics in Applied Mathematics. Academic Press, San Diego, CA, 1979.
- [41] M. V. Berry. Regular and irregular semiclassical wavefunctions. *Journal of Physics A: Mathematical and General*, **10**(12):2083–2091, Dec 1977.
- [42] M. V. Berry, M. Tabor, and J. M. Ziman. Level clustering in the regular spectrum. *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences*, **356**(1686):375–394, Sep 1977.
- [43] H. Bethe. Zur Theorie der Metalle. *Z. Physik*, **71**:205–226, Mar 1931.
- [44] G. D. Birkhoff. Proof of a recurrence theorem for strongly transitive systems. *Proceedings of the National Academy of Sciences*, **17**(12):650–655, Dec 1931.
- [45] G. D. Birkhoff. Proof of the ergodic theorem. *Proceedings of the National Academy of Sciences*, **17**(12):656–660, Dec 1931.
- [46] G. Biroli. A new kind of phase transition? *Nature Phys.*, **3**:222–223, Apr 2007.
- [47] O. Bohigas, M. J. Giannoni, and C. Schmit. Characterization of chaotic quantum spectra and universality of level fluctuation laws. *Phys. Rev. Lett.*, **52**(1):1–4, Jan 1984.
- [48] S. Boixo, T. Albash, F. M. Spedalieri, N. Chancellor, and D. A. Lidar. Experimental signature of programmable quantum annealing. *Nat. Commun.*, **4**:2067, Jun 2013.
- [49] S. Boixo, T. F. Rønnow, S. V. Isakov, Z. Wang, D. Wecker, D. A. Lidar, J. M. Martinis, and M. Troyer. Evidence for quantum annealing with more than one hundred qubits. *Nat. Phys.*, **10**(3):218–224, Feb 2014.
- [50] S. Boixo, V. N. Smelyanskiy, A. Shabani, S. V. Isakov, M. Dykman, V. S. Denchev, M. H. Amin, A. Y. Smirnov, M. Mohseni, and H. Neven. Computational multiqubit tunnelling in programmable quantum annealers. *Nat. Comm.*, **7**:10327, Jan 2016.
- [51] L. Boltzmann. Weitere Studien über das Wärmegleichgewicht unter Gasmolekülen. *Wiener Berichte*, **66**:275–370, 1982.
- [52] M. Boninsegni and S. Moroni. Population size bias in diffusion Monte Carlo. *Phys. Rev. E*, **86**(5):056712, Nov 2012.
- [53] M. Born and V. Fock. Beweis des Adiabatenatzes. *Z. Physik*, **51**:165–180, Mar 1928.
- [54] J. P. Bouchaud. Weak ergodicity breaking and aging in disordered-systems. *J. Phys. I France*, **2**(9):1705–1713, Sep 1992.

- [55] L. T. Brady and W. van Dam. Quantum Monte Carlo simulations of tunneling in quantum adiabatic optimization. *Phys. Rev. A*, **93**(3):032304, Mar 2016.
- [56] S. Bravyi, D. P. DiVincenzo, R. I. Oliveira, and B. M. Terhal. The complexity of stoquastic local hamiltonian problems. *Quant. Inf. Comp.*, **8**(5):361–385, May 2008.
- [57] S. Bravyi and B. Terhal. Complexity of stoquastic frustration-free hamiltonians. *SIAM J. Comput.*, **39**(4):1462–1485, 2010.
- [58] J. Bringewatt, W. Dorland, S. P. Jordan, and A. Mink. Diffusion Monte Carlo approach versus adiabatic computation for local Hamiltonians. *Phys. Rev. A*, **97**(2):022323, Feb 2018.
- [59] L. N. Bulaevskii, A. V. Zvarykina, Y. S. Karimov, R. B. Lyubovskii, and I. F. Shchegolev. Magnetic properties of linear conducting chains. *Zh. Eksp. Teor. Fiz.*, **62**(2):725, Aug 1972. [*Sov. Phys. JETP*, **35**(2):384, Aug 1972].
- [60] S. Butler, E. Coper, A. Li, K. Lorenzen, and Z. Schopick. Spectral properties of the exponential distance matrix. *arXiv e-prints*, Oct 2019. arXiv:1910.06373 [math.CO].
- [61] P. Calabrese and J. Cardy. Entanglement entropy and quantum field theory. *Journal of Statistical Mechanics: Theory and Experiment*, **2004**(06):P06002, Jun 2004.
- [62] G. Carleo and M. Troyer. Solving the quantum many-body problem with artificial neural networks. *Science*, **355**(6325):602–606, Feb 2017.
- [63] J. Carrasquilla, G. Torlai, R. G. Melko, and L. Aolita. Reconstructing quantum states with generative models. *Nat. Mach. Intell.*, **1**:155–161, Mar 2019.
- [64] C. E. Carroll and F. T. Hioe. Three-state systems driven by resonant optical pulses of different shapes. *J. Opt. Soc. Am. B*, **5**(6):1335–1340, Jun 1988.
- [65] D. M. Ceperley. Path integrals in the theory of condensed helium. *Rev. Mod. Phys.*, **67**(2):279–355, Apr 1995.
- [66] A. Chandran, I. H. Kim, G. Vidal, and D. A. Abanin. Constructing local integrals of motion in the many-body localized phase. *Phys. Rev. B*, **91**(8), Feb 2015.
- [67] K. Choo, T. Neupert, and G. Carleo. Two-dimensional frustrated $j_1 - j_2$ model studied with neural network quantum states. *Phys. Rev. B*, **100**(12):125124, Sep 2019.
- [68] D. T. Colbert and W. H. Miller. A novel discrete variable representation for quantum mechanical reactive scattering via the S-matrix Kohn method. *The Journal of Chemical Physics*, **96**(3):1982–1991, Feb 1992.

- [69] M. Collura, L. Dell’Anna, T. Felser, and S. Montangero. On the descriptive power of Neural-Networks as constrained Tensor Networks with exponentially large bond dimension. *arXiv e-prints*, May 2019. arXiv:1905.11351 [quant-ph].
- [70] J. F. Cornwell. *Group Theory in Physics: An Introduction*. Academic Press, Cambridge, 1997.
- [71] E. Crosson and A. W. Harrow. Simulated quantum annealing can be exponentially faster than classical simulated annealing. In *Ann. IEEE Symp. Found.*, pages 714–723. IEEE, Dec 2016.
- [72] J. E. Cuervo, P. N. Roy, and M. Boninsegni. Path integral ground state with a fourth-order propagator: Application to condensed helium. *J. Chem. Phys.*, **122**(11):114504, Mar 2005.
- [73] L. D’Alessio, Y. Kafri, A. Polkovnikov, and M. Rigol. From quantum chaos and eigenstate thermalization to statistical mechanics and thermodynamics. *Advances in Physics*, **65**(3):239–362, Jul 2016.
- [74] C. Dasgupta and S.-k. Ma. Low-temperature properties of the random Heisenberg antiferromagnetic chain. *Phys. Rev. B*, **22**(3):1305–1319, Aug 1980.
- [75] J. R. L. de Almeida and D. J. Thouless. Stability of the Sherrington–Kirkpatrick solution of a spin glass model. *Journal of Physics A: Mathematical and General*, **11**(5):983–990, May 1978.
- [76] A. De Luca, B. L. Altshuler, V. E. Kravtsov, and A. Scardicchio. Anderson localization on the Bethe lattice: Nonergodicity of extended states. *Phys. Rev. Lett.*, **113**(4):046806, Jul 2014.
- [77] A. De Luca and A. Scardicchio. Ergodicity breaking in a model showing many-body localization. *Europhys. Lett.*, **101**(3):37003, Feb 2013.
- [78] A. De Luca, A. Scardicchio, V. E. Kravtsov, and B. L. Altshuler. Support set of random wave-functions on the Bethe lattice. *arXiv e-prints*, Dec 2013. arXiv:1401.0019 [cond-mat.stat-mech].
- [79] A. de Touzalin, C. Marcus, F. Heijman, I. Cirac, R. Murray, and T. Calarco. Quantum manifesto for quantum technologies, Mar 2016.
- [80] V. S. Denchev, S. Boixo, S. V. Isakov, N. Ding, R. Babbush, V. Smelyanskiy, J. Martinis, and H. Neven. What is the computational value of finite-range tunneling? *Phys. Rev. X*, **6**(3):031015, Aug 2016.
- [81] B. Derrida. Random-energy model: An exactly solvable model of disordered systems. *Phys. Rev. B*, **24**(5):2613–2626, Sep 1981.

- [82] D. Deutsch. Quantum theory, the church–turing principle and the universal quantum computer. *Proceedings of the Royal Society of London. A*, **400**(1818):97–117, Jul 1985.
- [83] D. Deutsch and R. Jozsa. Rapid solution of problems by quantum computation. *Proceedings of the Royal Society of London. Series A*, **439**(1907):553–558, Dec 1992.
- [84] J. M. Deutsch. Quantum statistical mechanics in a closed system. *Phys. Rev. A*, **43**(4):2046–2049, Feb 1991.
- [85] V. Dotsenko. *An Introduction to the Theory of Spin Glasses and Neural Networks*. World Scientific Publishing Co., Inc., USA, 1994.
- [86] M. Droz. Spin glasses and the RKKY interaction. In A. Pękalski and J. Sznajd, editors, *Static Critical Phenomena in Inhomogeneous Systems*, pages 83–97, Berlin, Heidelberg, 1984. Springer Berlin Heidelberg.
- [87] F. J. Dyson. Statistical theory of the energy levels of complex systems. I. *Journal of Mathematical Physics*, **3**(1):140–156, Jan 1962.
- [88] J. T. Edwards and D. J. Thouless. Numerical studies of localization in disordered systems. *Journal of Physics C: Solid State Physics*, **5**(8):807–820, Apr 1972.
- [89] S. F. Edwards and P. W. Anderson. Theory of spin glasses. *Journal of Physics F: Metal Physics*, **5**(5):965–974, May 1975.
- [90] P. Ehrenfest and A. Ehrenfest. *The conceptual foundations of the statistical approach in mechanics*. Courier Corporation, 2015. [1959 reprint; transl. by M. J. Moravcskij].
- [91] A. Elgart and G. A. Hagedorn. A note on the switching adiabatic theorem. *Journal of Mathematical Physics*, **53**(10):102202, Sep 2012.
- [92] F. H. L. Essler, H. Frahm, F. Göhmann, A. Klümper, and V. E. Korepin. *The One-Dimensional Hubbard Model*. Cambridge University Press, Cambridge, 2005.
- [93] T. Esslinger. Fermi-Hubbard physics with atoms in an optical lattice. *Annual Review of Condensed Matter Physics*, **1**(1):129–152, 2010.
- [94] S. N. Evangelou and T. Ziman. The Anderson transition in two dimensions in the presence of spin-orbit coupling. *Journal of Physics C: Solid State Physics*, **20**(13):L235–L240, May 1987.
- [95] F. Evers and A. D. Mirlin. Anderson transitions. *Rev. Mod. Phys.*, **80**(4):1355–1417, Oct 2008.

- [96] D. Facchetti, P. Vivo, and G. Biroli. From non-ergodic eigenvectors to local resolvent statistics and back: A random matrix perspective. *EPL (Europhysics Letters)*, **115**(4):47003, Aug 2016.
- [97] K. Falconer. *Fractal Geometry: Mathematical Foundations and Applications*. Wiley, 2013.
- [98] L. Faoro, M. V. Feigel'man, and L. Ioffe. Non-ergodic extended phase of the quantum random energy model. *Annals of Physics*, **409**:167916, 2019.
- [99] E. Farhi, J. Goldstone, S. Gutmann, and D. Nagaj. How to make the quantum adiabatic algorithm fail. *Int. J. Quantum Inf.*, **6**(3):503–516, 2008.
- [100] E. Farhi, J. Goldstone, S. Gutmann, and M. Sipser. Quantum computation by adiabatic evolution. *arXiv e-prints*, Jan 2000. arXiv:quant-ph/0001106.
- [101] E. Farhi and S. Gutmann. Analog analogue of a digital quantum computation. *Phys. Rev. A*, **57**:2403–2406, Apr 1998.
- [102] A. L. Fetter and J. D. Walecka. *Quantum Theory of Many-particle Systems*. Dover Books on Physics. Dover Publications, Mineola, NY, 2003.
- [103] R. P. Feynman. Simulating physics with computers. *International Journal of Theoretical Physics*, **21**:467–488, Jun 1982.
- [104] A. B. Finnila, M. A. Gomez, C. Sebenik, C. Stenson, and J. D. Doll. Quantum annealing: a new method for minimizing multidimensional functions. *Chem. Phys. Lett.*, **219**(5–6):343–348, Mar 1994.
- [105] K. H. Fischer and J. A. Hertz. *Spin Glasses*. Cambridge Studies in Magnetism. Cambridge University Press, Cambridge, 1991.
- [106] D. S. Fisher. Random transverse field ising spin chains. *Phys. Rev. Lett.*, **69**(3):534–537, Jul 1992.
- [107] W. M. C. Foulkes, L. Mitas, R. J. Needs, and G. Rajagopal. Quantum Monte Carlo simulations of solids. *Rev. Mod. Phys.*, **73**(1):33–83, Jan 2001.
- [108] F. Franchini. *An Introduction to Integrable Techniques for One-Dimensional Quantum Systems*. Springer International, 1 edition, 2017.
- [109] J. Friedel. Metallic alloys. *Nuovo Cimento*, **7**:287–311, Sep 1958.
- [110] A. Garg. Tunnel splittings for one-dimensional potential wells revisited. *Amer. J. Phys.*, **68**:430, Apr 2000.
- [111] V. Gayrard. Aging in Metropolis dynamics of the REM: a proof. *arXiv e-prints*, Feb 2016. arXiv:1602.06081 [math.PR].

- [112] S. D. Geraedts, R. Nandkishore, and N. Regnault. Many-body localization and thermalization: Insights from the entanglement spectrum. *Phys. Rev. B*, **93**(17):174202, May 2016.
- [113] S. D. Geraedts, N. Regnault, and R. M. Nandkishore. Characterizing the many-body localization transition using the entanglement spectrum. *New Journal of Physics*, **19**(11):113021, Nov 2017.
- [114] S. Goldstein, J. L. Lebowitz, R. Tumulka, and N. Zanghì. Long-time behavior of macroscopic quantum systems. Commentary accompanying the English translation of John von Neumann’s 1929 article on the quantum ergodic theorem. *The European Physical Journal H*, **35**:173–200, Sep 2010.
- [115] I. V. Gornyi, A. D. Mirlin, and D. G. Polyakov. Interacting electrons in disordered wires: Anderson localization and low- T transport. *Phys. Rev. Lett.*, **95**(20):206603, Nov 2005.
- [116] L. K. Grover. A fast quantum mechanical algorithm for database search. *arXiv e-prints*, May 1996. arXiv:quant-ph/9605043.
- [117] A. K. Hartmann and M. Weigt. *Phase Transitions in Combinatorial Optimization Problems: Basics, Algorithms and Statistical Mechanics*. Wiley, Weinheim, 2005.
- [118] M. B. Hastings. An area law for one-dimensional quantum systems. *Journal of Statistical Mechanics: Theory and Experiment*, **2007**(08):P08024, Aug 2007.
- [119] M. B. Hastings. Obstructions to classically simulating the quantum adiabatic algorithm. *Quantum Inf. Comput.*, **13**(11–12):1038, Nov 2013. Appendix by M. H. Freedman.
- [120] B. Heim, T. F. Rønnow, S. V. Isakov, and M. Troyer. Quantum versus classical annealing of Ising spin glasses. *Science*, **348**(6231):215–217, Apr 2015.
- [121] W. Heisenberg. Zur Theorie des Ferromagnetismus. *Z. Physik*, **49**:619–636, 1928.
- [122] D. C. Herbert and R. Jones. Localized states in disordered systems. *Journal of Physics C: Solid State Physics*, **4**(10):1145–1161, Jul 1971.
- [123] J. Hubbard. Electron correlations in narrow energy bands. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, **276**(1365):238–257, 1963.
- [124] D. Hundertmark. A short introduction to Anderson localization. In P. Mörters, R. Moser, M. Penrose, M. Schwetlick, and J. Zimmer, editors, *Analysis and Stochastics of Growth Processes and Interface Models*, pages 194–218. Oxford University Press, Oxford, Sep 2008.

- [125] D. A. Huse, R. Nandkishore, and V. Oganesyan. Phenomenology of fully many-body-localized systems. *Phys. Rev. B*, **90**(17):174202, Nov 2014.
- [126] D. A. Huse, R. Nandkishore, V. Oganesyan, A. Pal, and S. L. Sondhi. Localization-protected quantum order. *Phys. Rev. B*, **88**(1):014206, Jul 2013.
- [127] T. Iadecola and M. Žnidarič. Exact localized and ballistic eigenstates in disordered chaotic spin ladders and the Fermi-Hubbard model. *Phys. Rev. Lett.*, **123**(3):036403, Jul 2019.
- [128] J. Z. Imbrie. On many-body localization for quantum spin chains. *Journal of Statistical Physics*, **163**(5):998–1048, 2016.
- [129] J. Z. Imbrie, V. Ros, and A. Scardicchio. Local integrals of motion in many-body localized systems. *Annalen der Physik*, **529**(7):1600278, May 2017.
- [130] E. M. Inack, G. Giudici, T. Parolini, G. Santoro, and S. Pilati. Understanding quantum tunneling using diffusion Monte Carlo simulations. *Phys. Rev. A*, **97**(3):032307, Mar 2018.
- [131] E. M. Inack and S. Pilati. Simulated quantum annealing of double-well and multiwell potentials. *Phys. Rev. E*, **92**(5):053304, Nov 2015.
- [132] E. M. Inack, G. E. Santoro, L. Dell’Anna, and S. Pilati. Projective quantum Monte Carlo simulations guided by unrestricted neural network states. *Phys. Rev. B*, **98**(23):235145, Dec 2018.
- [133] S. V. Isakov, G. Mazzola, V. N. Smelyanskiy, Z. Jiang, S. Boixo, H. Neven, and M. Troyer. Understanding quantum tunneling through quantum Monte Carlo simulations. *Phys. Rev. Lett.*, **117**(18):180402, Oct 2016.
- [134] W. Janke and T. Sauer. Properties of higher-order Trotter formulas. *Phys. Lett. A*, **165**(3):199–205, May 1992.
- [135] M. Jarret, S. P. Jordan, and B. Lackey. Adiabatic optimization versus diffusion Monte Carlo methods. *Phys. Rev. A*, **94**(4):042318, Oct 2016.
- [136] Z. Jiang, V. N. Smelyanskiy, S. Boixo, and H. Neven. Path-integral quantum Monte Carlo simulation with open-boundary conditions. *Phys. Rev. A*, **96**(4):042330, Oct 2017.
- [137] M. W. Johnson, M. H. S. Amin, S. Gildert, T. Lanting, F. Hamze, N. Dickson, R. Harris, A. J. Berkley, J. Johansson, P. Bunyk, E. M. Chapple, C. Enderud, J. P. Hilton, K. Karimi, E. Ladizinsky, N. Ladizinsky, T. Oh, I. Perminov, C. Rich, M. C. Thom, E. Tolkacheva, C. J. S. Truncik, S. Uchaikin, J. Wang, B. Wilson, and G. Rose. Quantum annealing with manufactured spins. *Nature*, **473**(7346):194–198, May 2011.

- [138] P. Jordan and E. Wigner. über das Paulische äquivalenzverbot. *Z. Physik*, **47**:631–651, Sep 1928.
- [139] R. Jördens, N. Strohmaier, K. Günter, H. Moritz, and T. Esslinger. A Mott insulator of fermionic atoms in an optical lattice. *Nature*, **455**(7210):204–207, Sep 2008.
- [140] A. Kamenev. *Field Theory of Non-Equilibrium Systems*. Cambridge University Press, New York, 2011.
- [141] T. Kasuya. A theory of metallic ferro- and antiferromagnetism on Zener’s model. *Progress of Theoretical Physics*, **16**(1):45–57, 07 1956.
- [142] K. Kechedzhi, V. Smelyanskiy, J. R. McClean, V. S. Denchev, M. Mohseni, S. Isakov, S. Boixo, B. Altshuler, and H. Neven. Efficient population transfer via non-ergodic extended states in quantum spin glass. *arXiv e-prints*, Jul 2018. arXiv:1807.04792 [quant-ph].
- [143] H. Kim and D. A. Huse. Ballistic spreading of entanglement in a diffusive nonintegrable system. *Phys. Rev. Lett.*, **111**(12):127205, Sep 2013.
- [144] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi. Optimization by simulated annealing. *Science*, **220**(4598):671–680, 1983.
- [145] J. A. Kjäll, J. H. Bardarson, and F. Pollmann. Many-body localization in a disordered quantum ising chain. *Phys. Rev. Lett.*, **113**(10):107204, Sep 2014.
- [146] S. Knysh. Zero-temperature quantum annealing bottlenecks in the spin-glass phase. *Nat. Commun.*, **7**:12370, Aug 2016.
- [147] K. Kobayashi, H. Sato, and M. Hoshi. The number of paths in boundary restricted Pascal triangle. In *2016 Information Theory and Applications Workshop (ITA)*, pages 1–10, 2016.
- [148] J. Kondo. Resistance minimum in dilute magnetic alloys. *Progress of Theoretical Physics*, **32**(1):37–49, Jul 1964.
- [149] V. E. Korepin. Universality of entropy scaling in one dimensional gapless models. *Phys. Rev. Lett.*, **92**(9):096402, Mar 2004.
- [150] B. Kramer and A. MacKinnon. Localization: theory and experiment. *Reports on Progress in Physics*, **56**(12):1469–1564, Dec 1993.
- [151] H. A. Kramers. Brownian motion in a field of force and the diffusion model of chemical reactions. *Physica*, **7**(4):284, Apr 1940.

- [152] V. Kravtsov, B. Altshuler, and L. Ioffe. Non-ergodic delocalized phase in Anderson model on Bethe lattice and regular graph. *Annals of Physics*, **389**:148—191, Feb 2018.
- [153] V. E. Kravtsov, I. M. Khaymovich, E. Cuevas, and M. Amini. A random matrix model with localization and ergodic transitions. *New Journal of Physics*, **17**(12):122002, Dec 2015.
- [154] L. D. Landau. О потерях энергии быстрыми частицами на ионизацию [On the energy loss of fast particles by ionization]. *USSR J. Phys*, **8**:201, 1944.
- [155] L. D. Landau and E. M. Lifshitz. *Quantum Mechanics: Non-Relativistic Theory*. Pergamon, New York, 3 edition, Dec 1981.
- [156] J. Langston. How the quest for a scalable quantum computer is helping fight cancer, Jul 2019.
- [157] T. Lanting, A. J. Przybysz, A. Y. Smirnov, F. M. Spedalieri, M. H. Amin, A. J. Berkley, R. Harris, F. Altomare, S. Boixo, P. Bunyk, N. Dickson, C. Enderud, J. P. Hilton, E. Hoskinson, M. W. Johnson, E. Ladizinsky, N. Ladizinsky, R. Neufeld, T. Oh, I. Perminov, C. Rich, M. C. Thom, E. Tolkacheva, S. Uchaikin, A. B. Wilson, and G. Rose. Entanglement in a quantum annealing processor. *Phys. Rev. X*, **4**(2):021041, May 2014.
- [158] C. R. Laumann, R. Moessner, A. Scardicchio, and S. L. Sondhi. Quantum annealing: the fastest route to quantum computation? *Eur. Phys. J. Special Topics*, **224**(1):75–88, Feb 2015.
- [159] C. R. Laumann, A. Pal, and A. Scardicchio. Many-body mobility edge in a mean-field quantum spin glass. *Phys. Rev. Lett.*, **113**(20):200405, Nov 2014.
- [160] C.-W. Liu, A. Polkovnikov, and A. W. Sandvik. Quasi-adiabatic quantum Monte Carlo algorithm for quantum evolution in imaginary time. *Phys. Rev. B*, **87**(17):174302, May 2013.
- [161] C.-W. Liu, A. Polkovnikov, and A. W. Sandvik. Quantum versus classical annealing: Insights from scaling theory and results for spin glasses on 3-regular graphs. *Phys. Rev. Lett.*, **114**(14):147203, Apr 2015.
- [162] D. J. Luitz, N. Laflorencie, and F. Alet. Many-body localization edge in the random-field Heisenberg chain. *Phys. Rev. B*, **91**(8):081103, Feb 2015.
- [163] P. Lunkenheimer, R. Wehn, U. Schneider, and A. Loidl. Glassy aging dynamics. *Phys. Rev. Lett.*, **95**(5):055702, Jul 2005.
- [164] S.-k. Ma, C. Dasgupta, and C.-k. Hu. Random antiferromagnetic chain. *Phys. Rev. Lett.*, **43**(19):1434–1437, Nov 1979.

- [165] G. D. Mahan. *Many Particle Physics*. Plenum Press, 3rd edition, 2000.
- [166] S. Mandrà, H. G. Katzgraber, and C. Thomas. The pitfalls of planar spin-glass benchmarks: raising the bar for quantum annealers (again). *Quantum Science and Technology*, **2**(3):038501, Jul 2017.
- [167] S. Mandrà, Z. Zhu, W. Wang, A. Perdomo-Ortiz, and H. G. Katzgraber. Strengths and weaknesses of weak-strong cluster problems: A detailed overview of state-of-the-art classical heuristics versus quantum approaches. *Phys. Rev. A*, **94**(2):022337, Aug 2016.
- [168] Y. I. Manin. *Вычислимое и Невычислимое [Computable and Uncomputable]*. Sovetskoye Radio, Moscow, 1980.
- [169] M. Marcovich (editor). *Hippolytus - Refutatio omnium haeresium*. Patristische Texte Und Studien (Band 25). De Gruyter, 1986.
- [170] R. Martoňák, G. E. Santoro, and E. Tosatti. Quantum annealing by the path-integral Monte Carlo method: The two-dimensional random Ising model. *Phys. Rev. B*, **66**(9):094203, Sep 2002.
- [171] J. C. Maxwell. II. Illustrations of the dynamical theory of gases. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, **20**(130):21–37, 1860.
- [172] J. C. Maxwell. V. Illustrations of the dynamical theory of gases. Part I. On the motions and collisions of perfectly elastic spheres. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, **19**(124):19–32, 1860.
- [173] J. C. Maxwell. On the dynamical theory of gases. *Philosophical Transactions of the Royal Society of London*, **157**:49–88, Jan 1867.
- [174] G. Mazzola, V. N. Smelyanskiy, and M. Troyer. Quantum Monte Carlo tunneling from quantum chemistry to quantum annealing. *Phys. Rev. B*, **96**(13):134305, Oct 2017.
- [175] G. B. Mbeng, L. Privitera, L. Arceci, and G. E. Santoro. Dynamics of simulated quantum annealing in random Ising chains. *Phys. Rev. B*, **99**(6):064201, Feb 2019.
- [176] C. C. McGeoch, R. Harris, S. P. Reinhardt, and P. I. Bunyk. Practical annealing-based quantum computing. *Computer*, **52**(6):38–46, Jun 2019.
- [177] M. L. Mehta. *Random Matrices*. Academic Press, 3rd edition, 2004.
- [178] R. G. Melko, G. Carleo, J. Carrasquilla, and J. I. Cirac. Restricted Boltzmann machines in quantum physics. *Nat. Phys.*, **15**:887–892, Jun 2019.

- [179] A. Messiah. *Quantum Mechanics*. Dover Books on Physics. Dover Publications, 1999.
- [180] N. Metropolis. The beginning of the Monte Carlo method. *Los Alamos Science (special issue dedicated to S. Ulam)*, pages 125–130, 1987.
- [181] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller. Equation of state calculations by fast computing machines. *The Journal of Chemical Physics*, **21**(6):1087–1092, Jun 1953.
- [182] M. Mezard and A. Montanari. *Information, Physics, and Computation*. Oxford University Press, Inc., Oxford, 2009.
- [183] M. Mézard and G. Parisi. Statistical physics of structural glasses. *Journal of Physics: Condensed Matter*, **12**(29):6655–6673, Jul 2000.
- [184] M. Mézard, G. Parisi, and M. A. Virasoro. *Spin Glass Theory and Beyond*. Lecture Notes in Physics Series. World Scientific, 1987.
- [185] Mézard, M., Parisi, G., Sourlas, N., Toulouse, G., and Virasoro, M. Replica symmetry breaking and the nature of the spin glass phase. *J. Phys. France*, **45**(5):843–854, May 1984.
- [186] F. Milde, R. A. Römer, M. Schreiber, and V. Uski. Critical properties of the metal-insulator transition in anisotropic systems. *The European Physical Journal B*, **15**(4):685–690, Jun 2000.
- [187] N. Minami. Local fluctuation of the spectrum of a multidimensional anderson tight binding model. *Comm. Math. Phys.*, **177**:709–725, Apr 1996.
- [188] C. C. Moore. Ergodic theorem, ergodic theory, and statistical mechanics. *Proceedings of the National Academy of Sciences*, **112**(7):1907–1911, Feb 2015.
- [189] G. Mossi, T. Parolini, S. Pilati, and A. Scardicchio. On the quantum spin glass transition on the Bethe lattice. *Journal of Statistical Mechanics: Theory and Experiment*, **2017**(1):013102, Jan 2017.
- [190] N. F. Mott. Conduction in glasses containing transition metal ions. *Journal of Non-Crystalline Solids*, **1**(1):1–17, 1968.
- [191] H. J. W. Müller-Kirsten. *Introduction to Quantum Mechanics: Schrödinger Equation and Path Integral*. World Scientific Publishing Co. Pte. Ltd., Singapore, 2006.
- [192] R. Nandkishore and D. A. Huse. Many-body localization and thermalization in quantum statistical mechanics. *Annual Review of Condensed Matter Physics*, **6**(1):15–38, 2015.

- [193] N. Nemec. Diffusion Monte Carlo: Exponential scaling of computational cost for large systems. *Phys. Rev. B*, **81**(3):035119, Jan 2010.
- [194] M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press, USA, 10th edition, 2011.
- [195] V. Oganesyan and D. A. Huse. Localization of interacting fermions at high temperature. *Phys. Rev. B*, **75**(15):155111, Apr 2007.
- [196] D. N. Page. Average entropy of a subsystem. *Phys. Rev. Lett.*, **71**(9):1291–1294, Aug 1993.
- [197] D. N. Page. Information in black hole radiation. *Phys. Rev. Lett.*, **71**(23):3743, Dec 1993.
- [198] R. G. Palmer. Broken ergodicity. *Advances in Physics*, **31**(6):669–735, 1982.
- [199] S. A. Parameswaran and R. Vasseur. Many-body localization, symmetry and topology. *Reports on Progress in Physics*, **81**(8):082501, Jul 2018.
- [200] G. Parisi. Infinite number of order parameters for spin-glasses. *Phys. Rev. Lett.*, **43**(23):1754–1756, Dec 1979.
- [201] G. Parisi. Mean field theory for spin glasses. *Physics Reports*, **67**(1):25–28, Dec 1980.
- [202] T. Parolini, E. M. Inack, G. Giudici, and S. Pilati. Tunneling in projective quantum Monte Carlo simulations with guiding wave functions. *Phys. Rev. B*, **100**(21):214303, Dec 2019.
- [203] T. Parolini and G. Mossi. Multifractal dynamics of the QREM. *arXiv e-prints*, Jul 2020. arXiv:2007.00315 [cond-mat.dis-nn].
- [204] D. Pekker, G. Refael, E. Altman, E. Demler, and V. Oganesyan. Hilbert-glass transition: New universality of temperature-tuned many-body dynamical quantum criticality. *Phys. Rev. X*, **4**(1):011052, Mar 2014.
- [205] P. Pfeuty. The one-dimensional Ising model with a transverse field. *Ann. Phys. (N. Y.)*, **57**(1):79–90, Mar 1970.
- [206] F. Pietracaprina, N. Macé, D. J. Luitz, and F. Alet. Shift-invert diagonalization of large many-body localizing spin chains. *SciPost Phys.*, **5**(5):45, Nov 2018.
- [207] S. Pilati, E. M. Inack, and P. Pieri. Self-learning projective quantum Monte Carlo simulations guided by restricted Boltzmann machines. *Phys. Rev. E*, **100**(4):043301, Oct 2019.

- [208] M. Pino, V. E. Kravtsov, B. L. Altshuler, and L. B. Ioffe. Multifractal metal in a disordered Josephson junctions array. *Phys. Rev. B*, **96**(21), Dec 2017.
- [209] A. C. Potter and R. Vasseur. Symmetry constraints on many-body localization. *Phys. Rev. B*, **94**(22):224206, Dec 2016.
- [210] P. Prelovšek, O. S. Barišič, and M. Žnidarič. Absence of full many-body localization in the disordered Hubbard chain. *Phys. Rev. B*, **94**(24):241104, Dec 2016.
- [211] J. Preskill. Quantum computing in the NISQ era and beyond. *Quantum*, **2**:79, Aug 2018.
- [212] I. V. Protopopov, W. W. Ho, and D. A. Abanin. Effect of SU(2) symmetry on many-body localization and thermalization. *Phys. Rev. B*, **96**(4):041122, Jul 2017.
- [213] I. V. Protopopov, R. K. Panda, T. Parolini, A. Scardicchio, E. Demler, and D. A. Abanin. Non-Abelian symmetries and disorder: A broad nonergodic regime and anomalous thermalization. *Phys. Rev. X*, **10**(1):011025, Feb 2020.
- [214] L. Reatto and G. L. Masserini. Shadow wave function for many-boson systems. *Phys. Rev. B*, **38**(7):4516–4522, Sep 1988.
- [215] M. Rigol, V. Dunjko, and M. Olshanii. Thermalization and its mechanism for generic isolated quantum systems. *Nature*, **452**:854–858, Apr 2008.
- [216] M. Rigol, V. Dunjko, V. Yurovsky, and M. 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**(5):050405, Feb 2007.
- [217] J. Roland and N. J. Cerf. Quantum search by local adiabatic evolution. *Phys. Rev. A*, **65**:042308, Mar 2002.
- [218] T. F. Rønnow, Z. Wang, J. Job, S. Boixo, S. V. Isakov, D. Wecker, J. M. Martinis, D. A. Lidar, and M. Troyer. Defining and detecting quantum speedup. *Science*, **345**(6195):420–424, Jul 2014.
- [219] V. Ros, M. Müller, and A. Scardicchio. Integrals of motion in the many-body localized phase. *Nuclear Physics B*, **891**:420–465, Feb 2015. [Corrigendum: *Nucl. Phys. B*, **900**:446–448, Nov 2015].
- [220] N. Rosenzweig and C. E. Porter. “repulsion of energy levels” in complex atomic spectra. *Phys. Rev.*, **120**(5):1698–1714, Dec 1960.

- [221] M. A. Ruderman and C. Kittel. Indirect exchange coupling of nuclear magnetic moments by conduction electrons. *Phys. Rev.*, **96**(1):99–102, Oct 1954.
- [222] A. Sarsa, K. E. Schmidt, and W. R. Magro. A path integral ground state method. *J. Chem. Phys.*, **113**(4):1366–1371, Jul 2000.
- [223] U. 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.
- [224] M. Schreiber, S. S. Hodgman, P. Bordia, H. P. Lüschen, M. H. Fischer, R. Vosk, E. Altman, U. Schneider, and I. Bloch. Observation of many-body localization of interacting fermions in a quasirandom optical lattice. *Science*, **349**(6250):842–845, Aug 2015.
- [225] A. Scollen and T. Hargraves. *Simulated Annealing: Introduction, Applications and Theory*. Mathematics research developments series. Nova Science Publishers, 2018.
- [226] S. Sen. Average entropy of a quantum subsystem. *Phys. Rev. Lett.*, **77**(1):1–3, Jul 1996.
- [227] M. Serbyn, Z. Papić, and D. A. Abanin. Local conservation laws and the structure of the many-body localized states. *Phys. Rev. Lett.*, **111**(12):127201, Sep 2013.
- [228] M. Serbyn, Z. Papić, and D. A. Abanin. Universal slow growth of entanglement in interacting strongly disordered systems. *Phys. Rev. Lett.*, **110**(26):260601, Jun 2013.
- [229] D. Sherrington and S. Kirkpatrick. Solvable model of a spin-glass. *Phys. Rev. Lett.*, **35**(26):1792–1796, Dec 1975.
- [230] P. W. Shor. Algorithms for quantum computation: discrete logarithms and factoring. In *Proceedings 35th Annual Symposium on Foundations of Computer Science*, pages 124–134, Santa Fe, NM, Nov 1994. IEEE.
- [231] P. W. Shor. Introduction to quantum algorithms. *arXiv e-prints*, Apr 2000. arXiv:quant-ph/0005003.
- [232] I. F. Silvera. The solid molecular hydrogens in the condensed phase: Fundamentals and static properties. *Rev. Mod. Phys.*, **52**(2):393–452, Apr 1980.
- [233] K. Slevin, P. Markoš, and T. Ohtsuki. Reconciling conductance fluctuations and the scaling theory of localization. *Phys. Rev. Lett.*, **86**(16):3594–3597, Apr 2001.
- [234] V. N. Smelyanskiy, K. Kechedzhi, S. Boixo, S. V. Isakov, H. Neven, and B. Altshuler. Nonergodic delocalized states for efficient population transfer within a narrow band of the energy landscape. *Phys. Rev. X*, **97**(1):011017, Jan 2020.

- [235] V. N. Smelyanskiy, K. Kechedzhi, S. Boixo, H. Neven, and B. Altshuler. Intermittency of dynamical phases in a quantum spin glass. *arXiv e-prints*, Jul 2019. arXiv:1907.01609 [cond-mat.dis-nn].
- [236] P. H. Song and D. L. Shepelyansky. Low-energy transition in spectral statistics of two-dimensional interacting fermions. *Phys. Rev. B*, **61**(23):15546–15549, Jun 2000.
- [237] S. Sorella and L. Capriotti. Green function Monte Carlo with stochastic reconfiguration: An effective remedy for the sign problem. *Phys. Rev. B*, **61**(4):2599–2612, Jan 2000.
- [238] M. Srednicki. Entropy and area. *Phys. Rev. Lett.*, **71**(5):666–669, Aug 1993.
- [239] M. Srednicki. Chaos and quantum thermalization. *Phys. Rev. E*, **50**(2):888–901, Aug 1994.
- [240] M. Srednicki. Thermal fluctuations in quantized chaotic systems. *Journal of Physics A: Mathematical and General*, **29**(4):L75–L79, Feb 1996.
- [241] M. Srednicki. The approach to thermal equilibrium in quantized chaotic systems. *Journal of Physics A: Mathematical and General*, **32**(7):1163–1175, Jan 1999.
- [242] L. Stella and G. E. Santoro. Quantum annealing of an Ising spin-glass by Green’s function Monte Carlo. *Phys. Rev. E*, **75**(3):036703, Mar 2007.
- [243] G. Stolz. An introduction to the mathematics of Anderson localization. In R. Sims and D. Ueltschi, editors, *Entropy and the Quantum II*, pages 71–108. Amer. Math. Soc., Providence, RI, 2011.
- [244] M. Suzuki. Transfer-matrix method and Monte Carlo simulation in quantum spin systems. *Phys. Rev. B*, **31**(5):2957–2965, Mar 1985.
- [245] E. Tarquini, G. Biroli, and M. Tarzia. Critical properties of the Anderson localization transition and the high-dimensional limit. *Phys. Rev. B*, **95**(9):094204, Mar 2017.
- [246] D. ter Haar. Foundations of statistical mechanics. *Rev. Mod. Phys.*, **27**(3):289–338, Jul 1955.
- [247] G. Theodorou and M. H. Cohen. Paramagnetic susceptibility of disordered N-Methyl-Phenazinium Tetracyanoquinodimethanide. *Phys. Rev. Lett.*, **37**(15):1014–1017, Oct 1976.
- [248] J. M. Thijssen. *Computational Physics*. Cambridge University Press, Cambridge, 2 edition, 2007.

- [249] C. K. Thomas and A. A. Middleton. Exact algorithm for sampling the two-dimensional Ising spin glass. *Phys. Rev. E*, **80**(4):046708, Oct 2009.
- [250] D. J. Thouless. A relation between the density of states and range of localization for one dimensional random systems. *Journal of Physics C: Solid State Physics*, **5**(1):77–81, Jan 1972.
- [251] D. J. Thouless, P. W. Anderson, and R. G. Palmer. Solution of ‘solvable model of a spin glass’. *The Philosophical Magazine: A Journal of Theoretical Experimental and Applied Physics*, **35**(3):593–601, 1977.
- [252] G. Toulouse. Theory of the frustration effect in spin glasses: I. In M. Mézard, G. Parisi, and M. Virasoro, editors, *Spin Glass Theory And Beyond: An Introduction To The Replica Method And Its Applications*, World Scientific Lecture Notes in Physics (Book 9), pages 99–103. World Scientific, Singapore, 1987. First published in *Communications on Physics* **2**:115–119 (1977).
- [253] C. J. Turner, A. A. Michailidis, D. A. Abanin, M. Serbyn, and Z. Papić. Weak ergodicity breaking from quantum many-body scars. *Nature Physics*, **14**(7):745–749, May 2018.
- [254] G. J. van den Berg. Anomalies in dilute metallic solutions of transition metals. In J. G. Daunt, D. O. Edwards, F. J. Milford, and M. Yaqub, editors, *Low Temperature Physics LT9*, pages 955–984, Boston, MA, 1965. Springer US.
- [255] J. Vannimenus and G. Toulouse. Theory of the frustration effect. II. Ising spins on a square lattice. *Journal of Physics C: Solid State Physics*, **10**(18):L537–L542, Sep 1977.
- [256] A. V. Vedyayev and V. A. Cherenkov. Low-temperature magnetic states in copper-manganese alloys. *Zh. Eksp. Teor. Fiz.*, **91**(6):2110–19, Dec 1986.
- [257] S. Vitiello, K. Runge, and M. H. Kalos. Variational calculations for solid and liquid ^4He with a “shadow” wave function. *Phys. Rev. Lett.*, **60**(19):1970, May 1988.
- [258] J. von Neumann. Beweis des Ergodensatzes und des H -Theorems in der neuen Mechanik. *Zeitschrift für Physik*, **57**:30–70, Jan 1929.
- [259] J. von Neumann. Proof of the quasi-ergodic hypothesis. *Proceedings of the National Academy of Sciences*, **18**(1):70–82, Jan 1932.
- [260] R. Vosk and E. Altman. Many-body localization in one dimension as a dynamical renormalization group fixed point. *Phys. Rev. Lett.*, **110**(6):067204, Feb 2013.
- [261] M. Žnidarič, T. Prosen, and P. Prelovšek. Many-body localization in the Heisenberg XXZ magnet in a random field. *Phys. Rev. B*, **77**(6):064426, Feb 2008.

- [262] U. Weiss, H. Grabert, P. Hänggi, and P. Riseborough. Incoherent tunneling in a double well. *Phys. Rev. B*, **35**(18):9535, Jun 1987.
- [263] E. Westerberg, A. Furusaki, M. Sigrist, and P. A. Lee. Low-energy fixed points of random quantum spin chains. *Phys. Rev. B*, **55**(18):12578–12593, May 1997.
- [264] S. R. White. Density matrix formulation for quantum renormalization groups. *Phys. Rev. Lett.*, **69**(19):2863–2866, Nov 1992.
- [265] E. Wigner. Characteristic vectors of bordered matrices with infinite dimensions. *Annals of Mathematics*, **62**(3):548–564, Nov 1955.
- [266] K. Yosida. Magnetic properties of Cu–Mn alloys. *Phys. Rev.*, **106**(5):893–898, Jun 1957.
- [267] A. P. Young, S. Knysh, and V. N. Smelyanskiy. Size dependence of the minimum excitation gap in the quantum adiabatic algorithm. *Phys. Rev. Lett.*, **101**:170503, Oct 2008.
- [268] J. Zakrzewski and D. Delande. Spin-charge separation and many-body localization. *Phys. Rev. B*, **98**(1):014203, Jul 2018.
- [269] Z. Zhu, A. J. Ochoa, and H. G. Katzgraber. Fair sampling of ground-state configurations of binary optimization problems. *Phys. Rev. E*, **99**(6):063314, Jun 2019.
- [270] M. Žnidarič, A. Scardicchio, and V. K. Varma. Diffusive and subdiffusive spin transport in the ergodic phase of a many-body localizable system. *Phys. Rev. Lett.*, **117**(4):040601, 2016.