



THÈSE DE DOCTORAT DE SORBONNE UNIVERSITÉ

réalisée au Laboratoire de Physique Théorique et Hautes Énergies

Sujet de la thèse :

Ergodicity Breaking in Disordered Matter: Novel Methods Across Physical Systems

présentée par **Greivin A. ALFARO-MIRANDA**

Présentée et soutenue publiquement à Paris le 10 Octobre 2025

devant le jury composé de :

Rapporteur: M. Juan J. Ruiz-Lorenzo Rapporteur: M. Dragi Karevski

Examinateur: M. Alexander K. Hartmann

Examinatrice: M^{me}. Jeanne Colbois Président: M. Nicolas Dupuis

Directrice de Thèse: M^{me}. Leticia F. Cugliandolo

Co-directeur de Thèse: M. Marco Tarzia

A Titi, mi abuelita, por haberme dado la sabiduría del rey Salomón

Résumé

L'hypothèse ergodique se situe au cœur de la mécanique statistique, mais de nombreux systèmes physiques s'en écartent lorsqu'ils sont hors d'équilibre. Les mécanismes à l'origine de cette brisure d'ergodicité donnent naissance à des phénomènes physiques riches, mais posent également des défis considérables, en particulier pour les simulations numériques. Au fil des années, diverses méthodes ont été développées pour relever ces défis dans des contextes spécifiques.

Cette thèse s'appuie sur ces approches établies et les adapte afin d'étudier des problèmes dans des systèmes physiques où de telles techniques ne sont pas encore courantes. Ce faisant, nous cherchons à mettre en évidence de nouvelles perspectives, tant sur les méthodes ellesmêmes que sur les systèmes désordonnés que nous examinons. Notre travail se concentre sur trois problèmes principaux dans l'étude plus large de la rupture d'ergodicité dans les systèmes désordonnés.

Premièrement, nous revisitons la correspondance entre les amas aléatoires (appelés amas FK–CK) et les modèles de spins sur réseau. Cette correspondance, fondée sur une probabilité de liaison, devient problématique dans les systèmes frustrés où cette probabilité peut devenir négative. En utilisant des techniques exactement solubles, telles que la méthode des cavités sur le réseau de Bethe, nous analysons la validité et les limites de la construction des amas FK–CK dans ces cas frustrés et proposons des extensions possibles lorsque l'ordre verre de spin émerge.

Deuxièmement, nous adaptons l'algorithme Monte Carlo SWAP — initialement développé pour les verres structuraux — aux modèles de spins sur réseau. Dans cette approche, les spins sont dotés de longueurs aléatoires, et des échanges de spins à longue portée sont alternés avec des mises à jour conventionnelles par retournement de spin unique, toutes deux acceptées selon le principe de l'équilibre détaillé. En appliquant cette méthode au modèle d'Edwards—Anderson bidimensionnel, nous montrons une accélération significative de la relaxation à basse température et un échantillonnage efficace de l'état fondamental avec un coût computationnel réduit. Cette adaptation apporte également un éclairage sur l'efficacité remarquable de la dynamique SWAP dans les systèmes de particules et met en évidence les liens entre l'évolution dynamique et les paysages d'énergie libre.

Enfin, nous étudions le rôle des résonances à l'échelle du système dans la transition de localisation à plusieurs corps (MBL) de la chaîne de spins XXZ en champ aléatoire. Nous utilisons un cadre récemment développé qui caractérise les statistiques des amplitudes de transmission entre des configurations multi-corps éloignées dans l'espace de Hilbert. Cette approche fournit un diagramme de phase en accord quantitatif avec d'autres méthodes récentes qui tiennent également compte des événements rares dans les observables en espace réel. De plus, nous proposons une image physique dans laquelle les propriétés de transport sont gouvernées par des chemins de délocalisation sur le graphe de l'espace de Hilbert (ou de Fock).

Pris ensemble, ces travaux illustrent l'intérêt de transférer des outils conceptuels et algorithmiques entre différents domaines, offrant des stratégies pratiques pour relever les divers défis posés par la rupture d'ergodicité dans les systèmes désordonnés.

Abstract

The ergodic hypothesis lies at the heart of statistical mechanics, yet many physical systems deviate from it under out-of-equilibrium conditions. The mechanisms behind this ergodicity breaking give rise to rich physical phenomena but also pose significant challenges, particularly for numerical simulations. Over the years, a variety of methods have been developed to address these challenges in specific contexts.

This thesis builds on these established approaches, adapting them to investigate problems in physical systems where such techniques are not yet standard. By doing so, we aim to uncover new insights into both the methods themselves and the disordered systems we are studying. Our work focuses on three main problems within the broader study of ergodicity breaking in disordered systems.

First, we revisit the mapping between random clusters (referred to as FK–CK clusters) and lattice spin models. This mapping, based on a bond probability, becomes problematic in frustrated systems where the probability can turn negative. Using exactly solvable techniques, such as the cavity method on the Bethe lattice, we analyze the validity and limitations of the FK–CK cluster construction in these frustrated cases and propose possible extensions when spin glass order emerges.

Second, we adapt the SWAP Monte Carlo algorithm—originally developed for structural glasses—to lattice spin models. In this approach, spins are endowed with random lengths, and long-range spin exchanges are alternated with conventional single-spin-flip updates, both accepted according to detailed balance. Applying this method to the two-dimensional Edwards—Anderson model, we demonstrate a significant acceleration of relaxation at low temperatures and efficient ground-state sampling with little computational cost. This adaptation also offers insights into why SWAP dynamics is so effective in particle systems and sheds light on connections between dynamical evolution and free-energy landscapes.

Finally, we investigate the role of system-wide resonances in the many-body localization (MBL) transition of the random-field XXZ spin chain. We employ a recently developed framework that characterizes the statistics of transmission amplitudes between distant many-body configurations in Hilbert space. This approach yields a phase diagram in quantitative agreement with other novel methods that also account for rare events in real-space observables. Moreover, we propose a physical picture in which transport properties are governed by delocalization paths on the Hilbert (or Fock) space graph.

Collectively, these studies illustrate the value of transferring conceptual and algorithmic tools across domains, yielding practical strategies for addressing the diverse challenges posed by ergodicity breaking in disordered systems.

Acknowledgments

The path to completing this thesis has been anything but easy. Technical hurdles were often accompanied by personal challenges, many of which arose from the experience of being an immigrant. Being far from home and building an often isolated life in an unfamiliar place brought its own frustrations, in addition to those inherent in pursuing an advanced technical degree. Yet alongside the undeniable privilege of this opportunity, I was fortunate to be accompanied by people I met along the way—friends and colleagues who helped me feel at home. Professionally, I had the privilege of working with extraordinary individuals whom I admire deeply.

First of all, I would like to thank my advisors, *Leticia Cugliandolo* and *Marco Tarzia*, for welcoming me as a PhD student and for the boundless patience I demanded at times. I am fortunate to look up to them not only as outstanding researchers and professionals, but also as remarkable human beings with whom I have had the privilege to spend time and from whom I have learned a great deal.

A Marco, che ringrazio per essere stato non solo il mio supervisore di dottorato, ma anche il mio insegnante personale d'italiano durante questi tre anni. Mi mancherà molto la costante incertezza su quale lingua avremmo parlato entrando nel tuo ufficio. Mille grazie per la tua inesauribile disponibilità; spero di non aver approfittato della tua gentilezza. Ti ringrazio anche per il tuo costante sostegno, persino nei momenti in cui dubitavo di me stesso più che mai. Il tuo incoraggiamento è stato per me una fonte inesauribile di motivazione, venendo da una persona che ammiro moltissimo come te.

A Leticia, a quien agradezco nunca haberme ahorcado durante el primer año, pero más importante aún: por haberme moldeado como físico. Espero que muchas de las cosas que aprendí con vos echen raíces para que se queden ahí para siempre. Muchas gracias por ser una excelente profesora durante la maestría y haberme motivado a seguir estudiando física estadística, gracias por ser un inspirador ¹ modelo a seguir y enseñarme que sí se puede hacer ciencia de donde venimos.

Apart from my supervisors—without whom this manuscript would not exist—I am also indebted to the members of the jury, Nicolas Dupuis, Jeanne Colbois, and Alexander Hartmann, and especially to the rapporteurs, Juan Jesús Ruiz-Lorenzo and Dragi Karevski, who agreed to read the manuscript. I am grateful to all of them for taking the time to listen, read and discuss my work.

A mi familia, especialmente mi mamá, mi tía Silma y mi abuelita, a quienes extraño profundamente y espero que sepan lo agradecido que estoy por haberme criado. Sin estas mujeres tan especiales en mi vida, no habría llegado muy lejos. Titi, gracias por aquellos años dejándome el almuerzo en la escuela para que yo siguiera estudiando. Mire que usted pensaba que se iba a morir antes de verme nacer; espero que se sienta tan feliz como yo al saber que no fue así y que pudo verme terminar un doctorado. A mi tía Silma y a mi mamá, a quienes les di mil y un dolores de cabeza, espero que este esfuerzo de concluir mis estudios se compare, aunque sea un poco, con el enorme esfuerzo que ustedes hicieron para darme todo lo que me dieron. Muchas gracias a mi mamá, por ser la mejor mamá y papá a la vez, y por tener el coraje y la valentía de sacarme adelante.

¹In this case I don't feel tempted to change the adjective to 'insightful'

A Dafne, que ao longo deste último ano foi meu alicerce emocional em um período especialmente extenuante. Esteve ao meu lado nos momentos de maior tensão enquanto eu escrevia este manuscrito e me fez sentir, pela primeira vez, estabilidade deste lado do mundo. Eu te amo e te admiro profundamente, e espero que assim permaneça pelo tempo que estiver ao nosso alcance.

Este viaje, desde mi perspectiva, comenzó en la maestría, al lado de la persona que más ha puesto a prueba mi paciencia y más veces me ha sacado de quicio, pero que al mismo tiempo terminó siendo una de las personas más fundamentales en mi vida: la bichilla primordial de Diana. Te amo muchísimo y no sé qué habría sido de mí en muchas ocasiones de no haber contado con tu existencia en mi vida.

A lo largo de estos cinco años, ha habido personas que llegaron a mi vida o se mantuvieron cerca de mí. A Uriel, a quien agradezco eternamente por ser una figura familiar y cálida, el mejor guía en un mundo ajeno. A Fabián, a quien admiro muchísimo como físico y ha sido un apoyo constante, con quien también puedo compartir la nostalgia por lo que dejamos atrás. Y a Adrián, quien más ha soportado mis peores actitudes y a quien aprecio profundamente, porque es un pequeño recordatorio de mi hogar y de esas cosas locales que extraño: Alajuela, la Capital del Mundo.

Agradezco a todos mis amigos que alguna vez me visitaron y rompieron mi rutina: Kuki, Jiménez, Eduardo, Niño, Manguz, Pepe y Luis; gracias por preservar nuestra amistad (o renovarla) a pesar de la distancia. También a los nuevos ticos que conocí al otro lado del mar y con quienes, milagrosamente, no cruzamos caminos en casa: gracias por construir un pequeño Costa Rica en París. A todos y cada uno de ellos, pero especialmente a Josué, Bryan, Cata y Andrea. I am also grateful to the other non-Costa Ricans I met, with whom I nevertheless felt equally understood: Fabián, Charbel, Bruna, and Gabriel.

Academically, I would also like to thank my comité de suivi de thèse, Benoît Estienne and Christophe Texier, for their support and timely check-ins. I am likewise grateful to the many professional physicists I have crossed paths with—whether through collaborations, such as Patrick Charbonneau, Giulio Biroli, Mingyuan Zheng, Nicolas Laflorencie, and Fabien Alet, or through courses and life, like Valentina Ros and Alberto Rosso—all of whom, in one way or another, helped me during these years.

Je souhaite également remercier Marco Picco et Michela Petrini pour leur soutien constant dans les questions administratives (et informatiques). Je suis aussi reconnaissant envers les autres membres du laboratoire avec qui j'ai eu l'occasion d'interagir — même brièvement — mais qui m'ont permis de me sentir partie intégrante de ce lieu, notamment Grégory Schehr, Benoît Estienne et Benoît Douçot. Agradezco particularmente a Redamy Pérez Ramos, a quien desafortunadamente llegué a conocer hasta el último año de mi doctorado. Le estoy profundamente agradecido por sus consejos profesionales y por su respaldo en mis decisiones hacia el futuro. Significó mucho para mí encontrar un hermano del Caribe en mi lugar de trabajo, aunque solo fuera durante el último año.

For my friends in the lab, starting with my original office mates who welcomed me warmly and gave me that very first sense of belonging: Pierre, Simon, Yehudi, Carlo (honorable mention for sharing his wonderful desk), and Jordan. To the younger generation as well—Rishi, Yehudi II: Rave Edition (Matteo) and the Fake Brit Tea Connoisseur (Arno). A very special thanks goes to my (temporary) late-night office mate, whom I deeply appreciate and respect. His warm company on those long nights when I stayed late in the lab made me feel far less alone: the man, the legend, Bing-Xin Lao. I am also grateful to my earliest interactions in the lab—the students from my year: Jules, Tom, Leonardo (traitor), Christiana, Sebastian, and the guy with the beard (the great writer Diyar with whom I never felt bored). You made the internship experience genuinely enjoyable and easygoing. Enfin, je tiens à remercier le paire fortement corrélée de physiciens statistiques du laboratoire, Mathis et Léo, qui

m'ont fait sentir que j'étais à ma place. Merci d'avoir été présents, de m'avoir toujours laissé l'espace pour parler et pour toutes ces innombrables discussions — souvent vagabondes et parfois sans issue — qui ont pourtant compté pour moi.

During these years I met some of the finest people I've encountered on this continent²: Andrei, Romullo, Francesco (1 and 2), Bianca, Andriani, Yann, Thorsten, Wendy, Federico (honorable mention for putting up with me while living together, with all the good and bad that entailed), Davide, Anthony, and many others that I may be forgetting (sorry). Thank you all for the convivial spaces—for drinking, talking, and watching movies—that I so badly needed. I'm especially grateful for the support you've given me, with special thanks to Andrei and Yann, with whom I felt most understood.

I also want to thank my adoptive family from this past year, with whom I had the pleasure of maintaining my early lunch schedule of 11:30 am and filled the year with coffee-time conversations—especially Dafne, Vincent, Sankarshan, Davide, Victor, and Elia, with whom I shared most of those moments.

Este proceso dista mucho de ser markoviano. Mi motivación y el deseo de realizar un doctorado en Física se remontan a mis años de bachillerato en Costa Rica. Aunque en este tiempo no he mantenido un contacto constante con mis compañeros de entonces, les debo buena parte del impulso que me llevó a tomar esta decisión: Anthony, Esteban, Luis, Andrés, Diego, Alejandro, Ari, Steven, Todd, Bryan y Fabián. A todas estas personas les agradezco el tiempo compartido en la EFis, donde comenzó el camino que me trajo hasta aquí y que, de alguna manera, encuentra su culminación en esta tesis.

²and also Vincent

Contents

In	trodu	action	2
1	The	Critical Clusters of Frustrated Spin Systems	32
	1.1	The spin configurations across the phase transition	34
	1.2	Geometrical percolation	
	1.3	The FK–CK correlated bond percolation	42
	1.4	The FK–CK clusters for frustrated models	47
	1.5	The RBIM on the Bethe lattice	50
	1.6	Percolation of the FK–CK clusters	53
	1.7	Spin glass clusters	61
	1.8	Conclusions and outlook	63
2	SWA	AP Dynamics for Frustrated Spin Systems	66
	2.1	Structural glasses and the SWAP method	68
	2.2	The 2d Edwards-Anderson (EA) model	
	2.3	The Δ -model for the clean Ising model	
	2.4	The frustrated Δ -model	
	2.5	Conclusions and outlook	
	A1	Branch-and-cut algorithm	106
	A2	More snapshots at $L=32$	
3 The	The	Importance of Rare Events in Many-Body Localization	108
	3.1	Anderson Localization	
	3.2	Many-body localization (MBL)	
	3.3	The Random Field XXZ Chain	
	3.4	Standard observables and their limitations	
	3.5	The large deviation method	
	3.6	Finite size analysis	155
	3.7	Rare vs typical samples	
	3.8	Conclusions and outlook	
	A1	Details of the numerical simulations	176
	A2	The computation of the error bars	177
	A3	The DPRM correlations	178
	A4	Supplementary results for the Anderson basis	
Co	nclu	sions and Perspectives	180

Introduction

From Atoms to Probabilities

The question of what the world is made of—whether it is continuous or composed of fundamental units—has fascinated human cultures for millennia. Across civilizations, thinkers have sought to explain the diversity of natural phenomena by imagining underlying constituents: tiny particles [1], 'seeds' [2, 3], or principles whose interactions give rise to the visible world. This search for the ultimate 'building blocks' of reality appears in many traditions: in India, the *Vaisheshika* school proposed minute eternal particles (*anu*) [4, 5]; early Buddhist and Jain philosophies envisioned aggregates of fundamental entities [6, 7]; Islamic scholars such as *Al-Ghazali* and the Ash'arite school developed atomistic theories to address change and creation [8, 9]; and Chinese natural philosophy, while often emphasizing continuous transformations through concepts like *yin—yang* and the *five phases*, also entertained notions of minimal constituents in Mohist thought [2, 3].

Modern physics, however, developed along a particular historical trajectory rooted in Europe, where the atomistic ideas of the Greek philosophers Democritus and Leucippus became especially influential [1]. Their vision of indivisible atoms moving in the void—combining and separating to form all things—endured through centuries and was later revitalized by the mechanistic worldview of Newtonian mechanics. In this framework, knowing the positions and velocities (or momenta) of all particles at a given instant, at least in principle, would allow one to predict the entire future of a physical system. This represents reductionism in its purest form: the belief that the microscopic laws, if followed faithfully, can explain everything about the macroscopic world [10, 11].

However, we encounter an insurmountable obstacle. Even the tiniest grain of sand contains on the order of 10^{19} atoms. Solving Newton's equations for each of them is impossible, not only practically but conceptually: the amount of information required is impossible to acquire—much less analyze—in a million lifetimes. There is, moreover, a deeper and subtler issue: even if we could follow every trajectory, the behavior of the whole may not always be reduced to the sum of its parts. Phenomena such as temperature, phase transitions, magnetism, or turbulence are emergent properties of matter: they arise only from the collective behavior of countless interacting constituents. Understanding this interplay between microscopic determinism and macroscopic emergence is at the heart of modern physics [12].

Faced with these challenges, physicists abandoned the purely deterministic perspective and adopted a probabilistic approach. Instead of solving every equation of motion, we describe the system statistically: macroscopic properties are understood as average quantities. The probabilities associated to the real physical observables do not reflect randomness in nature itself—the microscopic laws may still be deterministic—but rather our limited knowledge of the precise state of the system.

This shift, pioneered by Boltzmann and later formalized by Gibbs, was revolutionary. By focusing on probabilities rather than precise trajectories, statistical mechanics builds a bridge between the microscopic and macroscopic worlds. From this framework, we can predict measurable properties like temperature, pressure, magnetization, and even explain

dramatic collective phenomena like boiling or magnetic ordering (see [11, 13]).

These macroscopic quantities are the very observables we measure in the laboratory. Humans have been recording them for centuries: for example, Torricelli invented the mercury barometer in 1643 to measure atmospheric pressure; Galileo developed one of the first thermoscopes (a precursor to the thermometer) around 1593; and Fahrenheit introduced the mercury thermometer with a standardized scale in 1714. They underpin not only physics but also many vital domains of human activity, including climate science, energy production, and industrial processes.

Yet a subtle question remains: what is the meaning of the averages we calculate in theory? In experiments, we measure averages over time: the pressure of a gas on a wall, the temperature of a liquid, concentration of a solute in a mixture. These macroscopic quantities fluctuate due to the inaccessible microscopic dynamics underlying them and are therefore naturally described as outcomes of a stochastic process. What probability distribution, within our mathematical model, describes the fluctuations of these experimentally observed quantities?

The answer is provided by the *ergodic hypothesis*. It asserts that, for systems in equilibrium, the probability distribution describing long-time measurements of macroscopic quantities is the same as the distribution obtained by considering all microscopic states compatible with the macroscopic constraints of the system—such as total energy, temperature, or particle number. This equivalence allows us to compute macroscopic properties as *ensemble averages* from the statistical distribution of *microstates*, rather than by explicitly tracking their deterministic evolution.

In practice, this framework is tested not only through experiments on real physical systems, but also through numerical simulations of the underlying microscopic models. In the latter case, one simulates the mathematical model under specified conditions and compares the resulting observables with the predictions obtained from the ensemble description. Understanding this latter approach, and addressing the challenges it faces, constitutes the central focus of this thesis.

The Ergodic Hypothesis

Consider a system of classical particles—described by Newton's laws of motion, or any other equivalent mathematical formalism like Hamilton's equations. The state of a d-dimensional classical system composed of N particles is completely characterized by the canonical coordinate q_i , and its conjugate momentum p_i , associated to each individual particle (i = 1, ..., N).

As such, the relevant physical information of the system can be encoded in the Hamiltonian $\mathcal{H}=\mathcal{H}(\{\boldsymbol{p}_i\},\{\boldsymbol{q}_i\})$. In phase space Γ , the microstate of the system is defined by a 2Nd-dimensional vector $\boldsymbol{X}=(\{\boldsymbol{q}_i\},\{\boldsymbol{p}_i\})$ that indicates the position of a representative point in this high-dimensional space. The evolution of this representative point is fully determined by Hamilton's equations of motion; thus, in principle, describing the system amounts to solving these coupled deterministic equations for the N particles that compose it.

However, a macroscopic system typically contains on the order of 10^{23} particles. Even if one were able to write down a Hamiltonian that accurately describes the system, the complexity of its interactions and the presence of deterministic chaos would render such an approach practically futile. Instead, statistical mechanics provides a feasible way out of this conundrum by adopting the aforementioned probabilistic approach.

From a given initial position, the evolution of the representative point describes a trajec-

tory in phase space. We can estimate its position by defining a time-dependent probability distribution $\varrho(\boldsymbol{X};t)$ such that $\varrho(\boldsymbol{X};t)\,d\Gamma$ is the probability that the system is in within a volume $d\Gamma$ of the phase space around \boldsymbol{X} at time t. The time evolution of this probability distribution is given by

$$\frac{d\varrho}{dt} = \frac{\partial\varrho}{\partial t} + \sum_{i=1}^{Nd} \left(\frac{\partial\varrho}{\partial q_i} \frac{dq_i}{dt} + \frac{\partial\varrho}{\partial p_i} \frac{dp_i}{dt} \right) , \tag{1}$$

For a Hamiltonian system, the distribution function $\varrho(\boldsymbol{X};t)$ is constant along trajectories in phase space. Equivalently, the phase space volume is preserved under the time evolution generated by the Hamiltonian flow. This expressed by the Liouville equation

$$\frac{\partial \varrho}{\partial t} = -\sum_{i=1}^{Nd} \left(\frac{\partial \mathcal{H}}{\partial p_i} \frac{\partial \varrho}{\partial q_i} - \frac{\partial \mathcal{H}}{\partial q_i} \frac{\partial \varrho}{\partial p_i} \right) , \qquad (2)$$

where the expression on the right-hand side is the Poisson bracket $\{\mathcal{H},\varrho\}_{\mathrm{PB}}$

Without knowing the specific details involved in this evolution—that is a fundamental problem in statistical mechanics—one expects that in an equilibrium state of the system the probability distribution ϱ becomes stationary, yielding

$$\frac{\partial \varrho}{\partial t} = \{\mathcal{H}, \rho\}_{PB} = 0 \tag{3}$$

his stationary distribution $\varrho(\{q_i\}, \{p_i\})$ has a vanishing Poisson bracket with the Hamiltonian. Consequently, it can be expressed as a function of the conserved quantities of the system—most notably, the Hamiltonian itself—so that $\varrho = \varrho(\mathcal{H}(\{q_i\}, \{p_i\}))$.

Formally, the ergodic hypothesis states that, once the system has evolved for a sufficiently long time to reach equilibrium—characterized by the relaxation time τ_R —the following equality holds for any observable $\mathcal O$ measured over a time window T:

$$\overline{\mathcal{O}} \equiv \lim_{T \to \infty} \frac{1}{T} \int_{\tau_R}^{T + \tau_R} dt \, \mathcal{O}(\{\boldsymbol{q_i}(t)\}, \{\boldsymbol{p_i}(t)\}) = \int d\Gamma \, \varrho(\mathcal{H}(\{\boldsymbol{q_i}\}, \{\boldsymbol{p_i}\})) \, \mathcal{O}(\{\boldsymbol{q_i}\}, \{\boldsymbol{p_i}\})) \equiv \langle \mathcal{O} \rangle . \tag{4}$$

Although this statement cannot be formally proven for all physical systems (it has only been rigorously established for certain cases [14–18]), the predictive success of statistical mechanics has led to its widespread acceptance.

The Importance of Timescales

The validity of the equality in Eq. (4) removes the need to specify the detailed initial conditions of a system or to explicitly solve Hamilton's equations. However, many physical systems fail to satisfy the conditions required to fulfill the ergodic hypothesis. This phenomenon is referred to as *ergodicity breaking*.

One of the central questions in the study of ergodicity is determining what qualifies as a sufficiently long time in defining the relaxation time τ_R —the timescale over which a system reaches equilibrium, causing the ensemble probability distribution ϱ to become stationary. In many physical systems, τ_R increases with system size, $\tau_R(N)$, and can eventually exceed experimentally accessible timescales. Consequently, for macroscopic systems, τ_R may become extremely large, and observables measured over finite observation windows $\tau_{\mathcal{O}} \ll \tau_R$ may fail to reflect the true equilibrium value, instead capturing only transient, out-of-equilibrium behavior.

These apparent breakdowns of ergodicity are often associated with the emergence of new forms of order and, in many cases, with critical points where continuous (second-order) phase transitions occur and thermodynamic singularities appear. However, this connection is not universal. In particular, for structural glasses the nature of the transition remains debated: a clear order parameter is lacking, and it is still uncertain whether their dramatic dynamical slowdown corresponds to an underlying thermodynamic phase transition or a purely dynamical phenomenon [19–21].

While the loss of ergodicity signals the onset of rich and complex physical phenomena, it also poses major challenges for their theoretical and numerical study. Ferromagnets, spin glasses, and structural glasses provide paradigmatic examples of these mechanisms, illustrating both well-understood cases of spontaneous symmetry breaking and more subtle situations where the nature of ergodicity breaking in finite-dimensions is less understood.

Spontaneous Symmetry Breaking

In this thesis, we focus primarily on spin systems, which are simplified mathematical models of magnetic materials. The orientation of the atoms' magnetic moments is represented by a classical spin degree of freedom, $\sigma_i \in \{-1, +1\}$, assigned to each site of a d-dimensional lattice with N sites. For now, we will leave the geometry of the lattice to be generic.

The standard model in this framework is the Ising model, described by the Hamiltonian

$$\mathcal{H}(\{\sigma_i\}) = -J \sum_{\langle ij \rangle} \sigma_i \sigma_j - h_{\text{ext}} \sum_i \sigma_i , \qquad (5)$$

The two parameters in the model are the bond strengths J, that control the pairwise interaction between neighboring spins, and the external magnetic field $h_{\rm ext}$. For J>0, interactions are said to be ferromagnetic and they favor the alignment of spins, whereas J<0 defines anti-ferromagnetic bonds, and they favor the anti-alignment of neighboring spins where they point in opposite directions. The external field $h_{\rm ext}$ favors the alignment of the spins with its own direction (i.e. causing them to be aligned with $h_{\rm ext}$). We will consider the case of the ferromagnet where J>0.

The microstates of the system correspond to the 2^N possible spin configurations $\{\sigma_i\} = (\sigma_1, \sigma_2, \dots, \sigma_N) \in \mathscr{C}$, where \mathscr{C} denotes the configuration space. The thermodynamic properties of the system are often studied in the *canonical ensemble*, where the system is in contact with a thermal bath with inverse temperature³ $\beta = 1/T$. The equilibrium properties of the system are therefore dictated by the Gibbs-Boltzmann distribution

$$\varrho(\mathcal{H}(\{\sigma_i\}),\beta) = \frac{e^{-\beta\mathcal{H}(\{\sigma_i\})}}{\mathcal{Z}}, \qquad \mathcal{Z}(\beta) = \sum_{\{\sigma_i\}\in\mathscr{C}} e^{-\beta\mathcal{H}(\{\sigma_i\})}$$
 (6)

where $\mathcal{Z}(\beta)$ is the canonical partition function and the main observable is given by the global magnetization density, given by

$$m = \frac{1}{N} \sum_{i=1}^{N} \langle \sigma_i \rangle , \qquad (7)$$

with $\langle \cdots \rangle$ taken over the distribution of Eq. (6), and thus corresponds to an ensemble average.

In absence of the external field $h_{\text{ext}}=0$, the Hamiltonian has a discrete Z_2 symmetry as it is invariant under a global spin reversal $\mathcal{H}(\{\sigma_i\})=\mathcal{H}(\{-\sigma_i\})$. As a result, the ensemble average of σ_i vanishes i.e. $m=\langle\sigma_i\rangle=0$.

 $^{^3}$ We adopt natural units, setting $k_B=1$ throughout the thesis

One can verify experimentally that this is the case at high temperatures, where the material is demagnetized and any experimental average of its magnetic properties over time will yield a zero magnetization—apart from sub-extensive fluctuations—this defines the *paramagnetic phase*. However, upon lowering the temperature magnetic moments tend to align, and below some critical temperature T_c the system develops a spontaneous magnetization $m \neq 0$ signaling the emergence of a *ferromagnetic phase*.

This apparent contradiction is no contradiction at all, the system in the ferromagnetic phase has two possible equilibrium states related by symmetry: one with positive magnetization m>0, and another one with negative magnetization -m. These two states are fully degenerate, and thus compatible equilibrium states for the system at that temperature. For large enough times the system should visit these two states in the same proportion yielding m=0. However, this 'large enough' time is astronomically large and not accessible to the finite time windows $\tau_{\mathcal{O}}$ of experimental observations.

In practice, it is impossible to perfectly isolate a system from small external magnetic fields. As the experimentalist lowers the temperature through the transition, these residual fields bias the system toward one of the two symmetry-related orientations. In the theoretical model, the same effect can be described by introducing a small pinning field (or equivalently, boundary conditions of fixed spins pointing in the desired direction) that selects one of the two possible equilibrium states.

If the system is cooled below the critical temperature T_c with this small bias and then allowed to equilibrate, taking the thermodynamic limit $N\to\infty$ ensures that the system remains in the selected state. In this limit, even allowing $t\to\infty$ will not lead to nucleation events that restore symmetry: the system is locked into one configuration, and ergodicity is effectively broken. This is known as *spontaneous symmetry breaking*, as the selected equilibrium state breaks the Z_2 symmetry of the underlying Hamiltonian. Mathematically, is simply expressed as

$$\lim_{\delta h \to 0^{-}} \lim_{N \to \infty} \langle \sigma_i \rangle = -\lim_{\delta h \to 0^{+}} \lim_{N \to \infty} \langle \sigma_i \rangle \neq 0.$$
 (8)

This statement implies that the underlying phase space is partitioned in two disconnected components $\Gamma = \Gamma_+ \cup \Gamma_-$, where Γ_+ contains configurations with positive magnetization and Γ_- contains the symmetry related ones with negative magnetization. In other words, the Gibbs measure is split into sub-components, called *pure states*, and the ensemble average is taken as

$$\langle \cdots \rangle = \sum_{\alpha} w_{\alpha} \langle \cdots \rangle_{\alpha} . \tag{9}$$

For the ferromagnet there are two pure states, both of them related by symmetry and with equal probability $w_{\alpha} = 1/2$ (without any boundary conditions nor pinning fields), yielding

$$m = \langle \sigma_i \rangle = \frac{1}{2} \langle \sigma_i \rangle_+ + \frac{1}{2} \langle \sigma_i \rangle_- = 0.$$
 (10)

The fundamental thermodynamic properties of the microscopic model can be extracted from the free-energy density defined by

$$f(\beta) = -\frac{1}{\beta N} \ln \mathcal{Z}(\beta) . \tag{11}$$

In particular, the global minimum of the free energy density defines the equilibrium state in the thermodynamic limit $N \to \infty$.

Expressing the free-energy density as a function of the order parameter—in this case, the global magnetization density m—provides a clear illustration of spontaneous symmetry breaking. At high temperatures, the free-energy density exhibits a single global minimum at

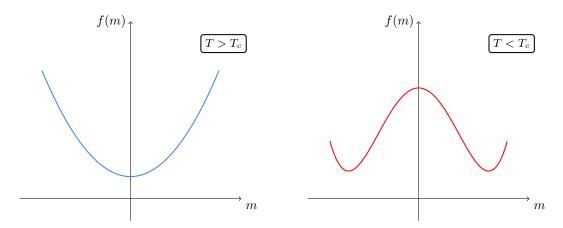


Fig. 1: Sketch of the free-energy density as a function of magnetization m. At high temperatures, $T > T_c$ (blue curve), the system has a single minimum at m = 0. At low temperatures, $T < T_c$ (red curve), two symmetric minima appear at nonzero m.

m=0, corresponding to a macroscopic equilibrium state—a unique pure state. However, as the temperature decreases, the free-energy density develops two symmetric minima at nonzero values of m, indicating the emergence of two pure states related by the symmetry $m \to -m$. This transition reflects the spontaneous breaking of the underlying symmetry, as shown in Fig. 1.

This mechanism of ergodicity breaking is straightforward. In a ferromagnet at low temperatures, the nature of the pure states is known a priori: they correspond to configurations with positive or negative magnetization. Consequently, we also know the form of the symmetry-breaking field that selects one of these states—an external field $h_{\rm ext}>0$ selects the positively magnetized state, while $h_{\rm ext}<0$ selects the negatively magnetized one.

While spontaneous symmetry breaking always implies ergodicity breaking, the reverse is not necessarily true. Ergodicity can break down through various mechanisms, including integrability, external driving, and quenched disorder—the latter being the primary focus of this thesis.

Disorder can be introduced into the Ising model by allowing both ferromagnetic $(J_{ij} > 0)$ and antiferromagnetic $(J_{ij} < 0)$ interactions between spins. At very low temperatures, it becomes impossible to orient all spins in a way that minimizes the energy of every bond simultaneously. In such cases, the system is said to be *frustrated*. An example is shown in Fig. 2, which illustrates a single frustrated plaquette in a square lattice. Here, the antiferromagnetic bond J_{41} remains unsatisfied, and flipping either of the interacting spins σ_1 or σ_4 merely shifts the frustration to a different bond.

At very low temperatures, the presence of quenched disorder can give rise to a new kind of order known as the spin glass phase, which exhibits a rich and complex phenomenology. In contrast to the ferromagnetic Ising model, where ergodicity breaking is understood as a consequence of spontaneous symmetry breaking, the mechanism of ergodicity breaking in spin glasses is more subtle.

In finite dimensions, models with a non-trivial spin glass phase transition are not exactly solvable, making it difficult to unambiguously characterize the nature of the breakdown of ergodicity. However, in the infinite-dimensional limit these systems are solvable, and the low-temperature spin glass phase reveals a qualitatively different form of ergodicity breaking. Specifically, the Gibbs measure decomposes into a multitude of pure states that are not related to one another by any simple symmetry, such as spin reversal.

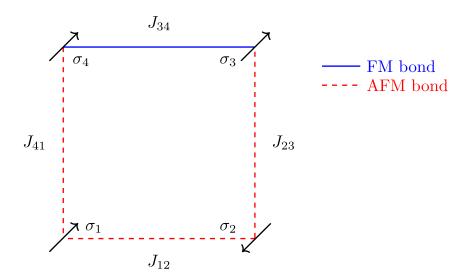


Fig. 2: Frustrated plaquette in a square lattice. A plaquette is frustrated when $\prod_{(ij)\in P} J_{ij} < 0$ where P is the set of all possible bonds, here $P = \{(12), (23), (34), (41)\}$. Ferromagnetic bonds (solid blue lines) favor aligned spins, while antiferromagnetic bonds (red dashed lines) are satisfied when spins are anti-aligned. The coexistence of both types of interactions leads to frustration in the plaquette, as exemplified by the unsatisfied bond J_{41} .

Glassiness

Spin glasses are metallic alloys composed of strongly interacting magnetic ions embedded in a weakly interacting, non-magnetic host material. These alloys are initially in a liquid state, and when rapidly cooled, the magnetic ions become fixed at random positions within the solid structure [22].

The interactions between pairs of magnetic ions are oscillatory and depend on their relative positions. These are described by the Ruderman–Kittel–Kasuya–Yosida (RKKY) interaction [23–25]

$$J(|\boldsymbol{r}_i - \boldsymbol{r}_j|) \sim \frac{\cos(\boldsymbol{k} \cdot (\boldsymbol{r}_i - \boldsymbol{r}_j))}{|\boldsymbol{r}_i - \boldsymbol{r}_j|^3},$$
(12)

where k is the wavevector, whose modulus is of the order of the Fermi vector.

The associated magnetic impurities—that is, the ions' magnetic moments—interact between each other through couplings that vary in both magnitude and sign. These interactions are sample-dependent, as each instance of the cooling process can fix the impurities at different random positions. These interactions between magnetic moments are thus a form of *quenched disorder*.

To capture the essential phenomenology of these systems, Edwards and Anderson [26] proposed a minimal model closely related to the Ising model of Eq. (2.7). This model consists of classical spins interacting on a lattice, but with random interactions reflecting the disorder introduced by the fixed, random positions of the ions.

In zero external magnetic field, the Hamiltonian is

$$\mathcal{H}_J(\sigma_i) = -\sum_{\langle ij\rangle} J_{ij}\sigma_i\sigma_j , \qquad (13)$$

The couplings J_{ij} are typically drawn from a random distribution, often taken to be either Gaussian,

$$\mathcal{P}(J_{ij}) = \frac{1}{\sqrt{2\pi J^2}} \exp\left(-\frac{J_{ij}^2}{2J^2}\right) , \qquad (14)$$

or bimodal,

$$\mathcal{P}(J_{ij}) = \frac{1}{2} \left[\delta(J_{ij} - J_0) + \delta(J_{ij} + J_0) \right]. \tag{15}$$

As a result, any given pair of spins can interact with ferromagnetic ($J_{ij} > 0$) or antiferromagnetic ($J_{ij} < 0$) bonds. This randomness introduces *frustration*: not all bonds can be simultaneously satisfied, and when frustration is strong enough, a spin glass phase emerges.

Unlike in a ferromagnet—where, at low temperatures, spins tend to align in the same direction—the equilibrium configurations of a spin glass at low temperatures consist of spins pointing in random directions. Despite this apparent disorder, most spins remain frozen in place, exhibiting only small thermal fluctuations, with only minor thermal fluctuations that do not lead to large-scale reorientation—unlike the paramagnetic phase.

The orientation of these frozen spins depend on the specific disorder realization of the bonds $\{J_{ij}\}$, which is why we explicitly label the Hamiltonian with the index J. The term 'glass' draws an analogy with structural glasses: just as window glass lacks the regular atomic structure of a crystal, spin glasses lack a clear magnetic ordering pattern. Nevertheless, even though samples are microscopically different, experimental spin glass realizations display the same average macroscopic behavior, indicative of a new type of thermodynamic order.

For a given disordered sample, the local orientation of the spins can be characterized by the local magnetization $m_i = \langle \sigma_i \rangle$. In the paramagnetic phase, this quantity vanishes because all spins fluctuate freely, yielding zero magnetization on average. In contrast, in the spin glass phase, spins remain frozen on average at each site, so m_i stays nonzero for a fixed realization of the disorder.

However, the sign of this frozen magnetization can vary between different disorder realizations. To eliminate this sign ambiguity, we consider m_i^2 and then average over many disorder realizations. This leads to the definition of the global order parameter

$$q_{\rm EA} \equiv \frac{1}{N} \sum_{i=1} [m_i^2] = \frac{1}{N} \sum_{i=1} [\langle \sigma_i \rangle^2] ,$$
 (16)

where $[\cdots]$ denotes the average over disorder realizations $\{J_{ij}\}$. This is the Edwards–Anderson (EA) order parameter, originally introduced in the seminal paper describing the model in Eq. (13) [26].

In principle, we could extract this information from the calculation of the free-energy density of the model. However, the presence of quenched disorder results in a sample dependent free-energy density $f_J(\beta) = \frac{1}{\beta N} \ln \mathcal{Z}_J(\beta)$. The relevant quantity to be probed is instead the *quenched averaged* free-energy density

$$f = [f_J(\beta)] = -\frac{1}{\beta N} [\ln \mathcal{Z}_J(\beta)] . \tag{17}$$

In finite dimensions, an exact analytical computation of the free-energy density is not feasible, and one must rely on approximations—most notably the *mean-field approximation*, which becomes exact in the limit of infinite spatial dimensions $(d \to \infty)$. In this limit, the system loses its short-range character, and the underlying lattice becomes equivalent to the fully-connected case, meaning that each spin interacts with every other spin in the system. The Hamiltonian in this mean-field limit takes the form

$$\mathcal{H}_J(\{\sigma_i\}) = -\frac{1}{2} \sum_{i \neq j} J_{ij} \sigma_i \sigma_j , \qquad (18)$$

where the factor of 1/2 avoids double-counting spin pairs. To ensure that the free-energy remains extensive in the thermodynamic limit, the variance of the coupling distribution $\mathcal{P}(J_{ij})$ is rescaled by the system size N.

This model is known as the Sherrington-Kirkpatrick model [27, 28]. In this context, the average of the quenched free-energy density $[f_J(\beta)]$ can be obtained by the *replica trick*

$$[\ln \mathcal{Z}_J] = \lim_{n \to 0} \frac{[\mathcal{Z}_J^n] - 1}{n} . \tag{19}$$

The core idea is to compute \mathcal{Z}_J^n for integer n and then perform an analytical continuation to calculate the limit $n \to 0$. In this middle step, we define a *replicated partition function*

$$\mathcal{Z}_{J}^{n} = \sum_{\{\sigma_{i}^{(a)}\}} \exp\left(-\beta \sum_{a} \mathcal{H}_{J}\left(\left\{\sigma_{i}^{(a)}\right\}\right)\right) , \qquad (20)$$

where we have introduced a new index a to distinguish between spin degrees of freedom corresponding to the n different replicas.

These replicas serve as a formal tool for computing the quenched average of the free-energy density. They are introduced as n non-interacting copies of the system, all sharing the same disorder realization $\{J_{ij}\}$. In the calculation of the free energy—details omitted here—averaging over the disorder decouples the site degrees of freedom, but at the cost of introducing effective correlations between replicas.

From this calculation, after performing a saddle-point calculation, we identify the order parameter is identified as the overlap between two replicas a and b

$$q_{ab} \equiv \frac{1}{N} \sum_{i=1}^{N} \sigma_i^{(a)} \sigma_i^{(b)} ,$$
 (21)

but we still need to take the subtle limit of $n \to 0$, where n is the dimension of the overlap matrix q_{ab} . To do this we need to find a parametrization of the matrix q_{ab} as a function its entries and dimension n.

In principle, nothing distinguishes one replica from another, as they are introduced purely as a mathematical artifact to aid the calculation. This observation led Sherrington and Kirkpatrick to propose the *replica symmetric* (RS) solution, which assumes

$$q_{ab} = q_0 + (1 - q_0)\delta_{ab} , (22)$$

meaning that all distinct replicas have the same mutual overlap q_0 , while the overlap of a replica with itself is exactly one, since it corresponds to the same spin configuration.

Evaluating the limit yields the extrema of the free-energy, which admit two possible solutions: $q_0 = 0$, corresponding to the paramagnetic phase, and $q_0 \neq 0$, identifying the spin glass phase below a finite transition temperature $T_{\rm SG}$. However, this replica symmetric solution presents a serious issue: the entropy becomes negative at zero temperature. Furthermore, de Almeida and Thouless proved that the saddle-point corresponding to the RS solution is unstable in the spin glass phase. Consequently, the replica symmetric ansatz must be replaced by a more general formulation to correctly describe the spin glass phase.

The correct solution needed the breaking of the replica symmetry (i.e. $q_{ab} \neq q_0$ for $a \neq b$) meaning that different replicas are non-trivially correlated. The correct replica symmetry breaking solution was introduced and described by Parisi in a list of seminal articles [29–34]. The basic idea is that between any pair of replicas, they can take different k+1 ($\leq n$) overlap values. This is known as k-step replica symmetry breaking (RSB). The value of these replicas are ordered in a hierarchical structure [35]

$$0 \le q_0 \le q_1 \le \dots \le q_{K-1} \le q_k = 1. \tag{23}$$

and the overlap matrix q_{ab} is arranged in black sub-matrices of size m_j according to these values where

$$q_{ab} = q_i \quad \text{if} \quad m_{i+1} \le |a - b| \le m_i \,, \tag{24}$$

and the sub-matrix sizes are organized as

$$n = m_0 \ge m_1 \ge \dots \ge m_k \ge m_{k-1}$$
 satisfying $\frac{m_{j+1}}{m_j} \in \mathbb{N}$. (25)

a representation of this construction is shown in Fig. 3. For 1-step RSB the off diagonal elements of the matrix can have two possible overlap values q_0 and q_1 . In general, for k-step RSB we will have k+1 possible values of the off-diagonal elements.

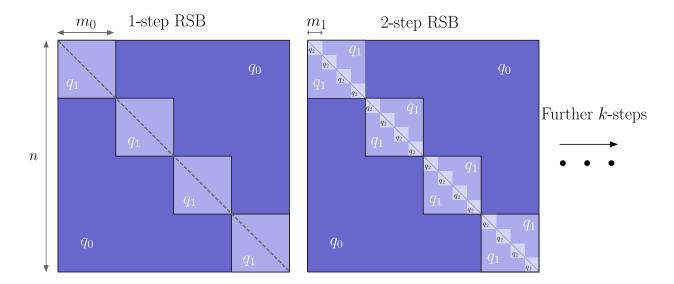


Fig. 3: Construction of the overlap matrix following Parisi's hierarchical scheme. The main diagonal, shown as gray dashed lines, is fixed to 1, representing the overlap of a configuration with itself. For k = 1, the matrix is divided into large sub-matrices whose off-diagonal elements are all equal to q_1 . At k = 2, each of these sub-matrices is further subdivided into smaller blocks, with the off-diagonal elements inside each new block set to q_2 . Image adapted from Ref. [36].

The correct SK solution corresponds to the case of full replica symmetry breaking (f-RSB) where the limit $k \to \infty$ is taken. In this limit, we obtain a physical value for the entropy at low temperatures, recovering the expected result of being zero at zero temperature. This solution has been recently proven to be exact for mean-field models by two mathematical physics, F. Guerra [37] and M. Talagrand [38].

This solution is not merely a mathematical device to reach the correct result—it encapsulates essential physical features of the spin glass phase in the mean-field limit. In this phase, the off-diagonal q_{ab} represent the possible values of the overlap between different equilibrium configurations. Recalling the decomposition of the Gibbs measure into pure state components α , this indicates that, in the spin glass phase, the phase space itself is fragmented into a multitude of pure states. These states are related to each other through their mutual overlaps.

The set of possible overlaps between pure states depends on the level of RSB used to describe the low-temperature phase. For instance, in the case of 1-step RSB, the pure states are grouped into two distinct clusters: within the same cluster, all states share the same overlap q_0 , whereas states belonging to different clusters have a mutual overlap q_1 . As k increases, the structure of overlaps becomes more intricate.

In this context, recall that the thermal average of any observable can be expressed as a sum over the contributions from the different pure states—partitioning the phase space into sub-components of the Gibbs measure

$$\langle \cdots \rangle = \sum_{\alpha} w_{\alpha}^{J} \langle \cdots \rangle_{\alpha} , \qquad (26)$$

where $\langle \cdots \rangle_{\alpha}$ denotes the thermal average restricted to the pure state α , and w_{α}^{J} is the statistical weight of that state for a given realization of the quenched disorder J, with $\sum_{\alpha} w_{\alpha}^{J} = 1$. In the spin glass phase, unlike in the ferromagnet, these weights fluctuate from sample to sample.

We can then introduce the overlap between two pure states α and β

$$q_{\alpha\beta}^{J} \equiv \frac{1}{N} \sum_{i=1}^{N} \langle \sigma_i \rangle_{\alpha} \langle \sigma_i \rangle_{\beta} = \frac{1}{N} \sum_{i=1}^{N} m_i^{(\alpha)} m_i^{(\beta)} , \qquad (27)$$

where $\langle \sigma_i \rangle_{\alpha}$ denotes the thermal average of spin *i* restricted to the pure state α . In other words, the overlap of local magnetizations between two pure states α and β .

If we consider a single pure state $\alpha=\beta$ and average over disorder realizations, we recover the definition of the Edwards–Anderson (EA) order parameter. In this sense, $q_{\rm EA}$ characterizes the properties of a single pure state α , a description that is exact at the transition temperature $T_{\rm SG}$. However, as the temperature decreases below $T_{\rm SG}$, additional pure states appear, and the EA order parameter becomes blind to this richer structure of the phase space.

Instead, one needs to look into the statistics of different overlaps along the low-temperature phase, that is characterized by the probability distribution

$$\mathcal{P}_J(q) = \sum_{\alpha\beta} w_{\alpha}^J w_{\beta}^J \delta(q - q_{\alpha\beta}^J) . \tag{28}$$

Consequently, $q_{\rm EA}$ is directly related to the self-overlap $q_{\alpha\alpha}$, that corresponds to the maximum overlap value.

Therefore, the most complete order parameter for the spin glass transition is the *overlap* probability distribution $\mathcal{P}(q)$, defined as

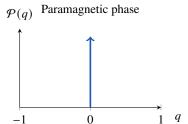
$$\mathcal{P}(q) = [\mathcal{P}_J(q)] = \sum_{\alpha,\beta} w_{\alpha}^J w_{\beta}^J \delta(q - q_{\alpha\beta}^J), \qquad (29)$$

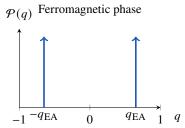
This quantity clearly distinguishes between different phases:

- In the paramagnetic phase, $\mathcal{P}(q) = \delta(q)$, since typical overlaps vanish.
- In the ferromagnetic phase (are any other replica-symmetric phase, e.g. the spin glass exactly at $T_{\rm SG}$) $\mathcal{P}(q)=\frac{1}{2}\delta(q-q_{\rm EA})+\frac{1}{2}\delta(q+q_{\rm EA})$, where $q_{\rm EA}=m^2$ for the ferromagnet.
- In a k-step RSB phase, $\mathcal{P}(q)$ contains multiple delta peaks corresponding to distinct overlap values between pure states. As $k \to \infty$, in the full RSB (f-RSB) phase, these peaks merge into a continuous distribution, and the possible overlaps form an interval, as illustrated in Fig. 4.

In this sense, the entire low-temperature phase $T < T_{\rm SG}$ in the SK model is 'critical' as more pure states appear upon decreasing the temperature—these pure states are organized in a ultrametric structure [39], with a distance between them defined by their overlaps

$$d(\alpha, \beta) = \frac{1}{2}(q_{\text{EA}} - q_{\alpha\beta}) . \tag{30}$$





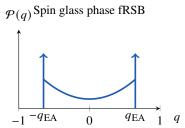


Fig. 4: Different forms of the overlap distribution $\mathcal{P}(q)$ for three different phases.

Translating the free-energy plots of Fig. 1 to the mean-field spin-glass setting is subtle. There is no single global magnetization; instead the system is described by the N local magnetizations $\{m_i\}$, which span a high-dimensional free-energy landscape. This landscape contains many global minima, each corresponding to the set of local magnetizations of a pure state α , denoted $\{m_i^{(\alpha)}\}$. They are separated by barriers whose heights grow with N and become effectively infinite in the thermodynamic limit. This leads to ergodicity breaking: once the system is in a given pure state, equilibrium dynamics cannot take it to another on any finite timescale.

Beyond pure states, the free-energy landscape also exhibits an exponential number of higher-free-energy local minima (*metastable states*). These do not dominate the equilibrium measure but govern the out-of-equilibrium relaxation from random initial conditions. These metastable states appear from local minima of the Thouless-Anderson-Palmer (TAP) free-energy [40]. This TAP free-energy is calculated from self-consistent equations of the local magnetizations, that were introduced to describe the spin glass phase before Parisi's RSB solution. The presence of such a vast number of metastable states makes the dynamics extremely slow: the system can wander among many nearby minima—or even flat regions—and become trapped for long times before finding deeper valleys [41–43].

Some schematic representations of these free-energy landscapes are shown in Fig. 5. These are three-dimensional projections of an inherently high-dimensional space, so they should be interpreted with caution—important features, such as the proliferation of nearly flat regions, are not explicitly visible.

Conceptually, these sketches illustrate the mechanism of ergodicity breaking in mean-field spin glasses. Unlike the ferromagnetic case shown in Fig. 1, where ergodicity is broken between just two symmetry-related pure states, spin glasses possess a vast multiplicity of pure states along with numerous long-lived metastable states. Dynamically, the system can become trapped in these states for extremely long times, preventing full exploration of phase space within accessible observation windows.

The phenomenology described above for the spin-glass phase is strictly valid only above the upper critical dimension $D_u = 6$ [45]. Whether Parisi's mean-field picture persists in lower dimensions remains an open and debated question [46], and several alternative theoretical descriptions have been proposed for finite-dimensional systems.

One of the main competing approaches is the droplet theory of spin glasses [47]. This framework proposes that the low-temperature phase contains only two equilibrium states, related to each other by a global spin-flip symmetry. In this view, the mechanism of ergodicity breaking is essentially the same as in a standard ferromagnet. The crucial difference is that here the pure states are composed of spins pointing in many different directions—determined by the specific realization of the quenched disorder—rather than in a uniform orientation. Consequently, the corresponding symmetry-breaking field is a random field, not a simple uniform one. In this framework, the spin glass can be thought of as a 'disguised ferromagnet'

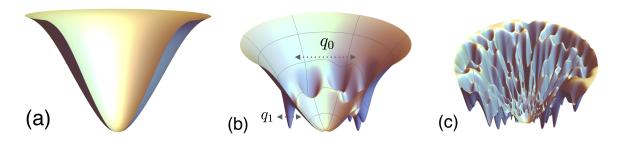


Fig. 5: Sketches of the free-energy density landscape close to equilibrium. (a) A single minimum, typical of high temperatures (paramagnetic phase). (b) Multiple minima in a 1-step RSB model at temperatures below $T_{\rm SG}$. The parameters q_0 and $q_1=q_{\rm EA}$ indicate the overlap between configurations of different equilibrium states and configurations within the same one, respectively. (c) The hierarchical organization of equilibrium states with relatively low barriers between them as realized in the Sherrington-Kirkpatrick model. Image taken from Ref. [44].

exhibiting conventional spontaneous symmetry breaking but in a highly nontrivial, disorder-dependent form.

The Dynamical Point of View

In essence, ergodicity breaking is a statement about the dynamics of a system and can often be identified by the system's persistent memory of its initial conditions. In what follows, we introduce key observables commonly used in numerical and theoretical studies to probe this dynamical behavior.

Unlike the position and momenta of classical particles, Ising degrees of freedom are bimodal variables with no well defined time derivative and have no intrinsic dynamics—also seen from the vanishing of its Poisson bracket $\{\sigma_i, \mathcal{H}\}$. To associate a dynamical evolution to the spin degrees of freedom, one assumes that the presence of a thermal bath induces fluctuations in the form of stochastic spin moves.

When a system is in contact with a thermal bath at temperature T, the coupling with the bath tends to drive it towards thermal equilibrium at that temperature. In equilibrium, the microscopic configurations follow a stationary probability distribution and satisfy microscopic reversibility: the probability of transitioning from one equilibrium configuration $\mathcal C$ to another $\mathcal C'$ is exactly equal to that of the reverse process.

This condition is known as detailed balance, which for spin configurations reads

$$\mathbb{P}_{eq}(\{\sigma_i\})\mathbb{P}(\{\sigma_i\} \to \{\sigma_i'\}) = \mathbb{P}_{eq}(\{\sigma_i'\})\mathbb{P}(\{\sigma_i'\} \to \{\sigma_i\}), \qquad (31)$$

where $\mathbb{P}_{eq}(\{\sigma_i\})$ is the equilibrium probability of finding the system in configuration σ_i —for a spin system in equilibrium, this is given by the Gibbs–Boltzmann distribution of Eq. (6)—and $\mathbb{P}(\sigma_i \to \sigma_i')$ is the transition probability from configuration σ_i to σ_i' .

In practice, these transition probabilities are often implemented via the Metropolis–Hastings algorithm, that accepts or rejects possible microscopic moves according to the probability.

$$\mathbb{P}_{\rm acc} = \min(1, e^{-\beta \Delta E}) \tag{32}$$

where the energy difference ΔE is the difference between the current configuration $\{\sigma_i\}$ and the proposed configuration $\{\sigma_i\}$ for the transition $\sigma_i \to \sigma'_i$.

Although many choices of microscopic moves are possible [48], in the present discussion we restrict ourselves to the standard single-spin-flip dynamics, in which the proposed move consists solely of flipping the spin at a randomly chosen site $\sigma_i \to -\sigma_i$.

In this setting, we investigate how the system approaches equilibrium through a *thermal quench*. In such a process, the system is first prepared in a configuration representative of a given temperature $T_{\rm initial}$. Then, this configuration is suddenly 'quenched' to a new temperature $T_{\rm final}$. The term *quench* refers to an abrupt change in a control parameter (here, the temperature), after which the system evolves according to the chosen dynamics. At the new temperature, the microscopic configurations change step by step, and—over time—the system relaxes toward the equilibrium state corresponding to $T_{\rm final}$.

In the presence of a second-order phase transition, the dynamical behavior of the configurations strongly depends on the relative values of the initial and final temperatures with respect to the critical temperature. The spatial correlations of the configurations can differ substantially between the distinct phases, and the relaxation dynamics may become very slow. To illustrate these effects, we focus once again on the ferromagnetic Ising model.

The Ferromagnetic Ising model

In thermal equilibrium, approaching the critical temperature T_c from the high-temperature phase, the equilibrium correlation length $\xi_{\rm eq}$ diverges. This signals the emergence of long-range spatial order, as spins become correlated over arbitrarily large distances. Near the transition, the divergence follows the scaling law

$$\xi_{\rm eq} \simeq |T - T_c|^{-\nu} \,, \tag{33}$$

where ν is a universal critical exponent.

The initial configuration for a quench is necessarily representative of some equilibrium phase. A perfectly aligned state corresponds to $T_{\rm initial}=0$, while a completely random spin configuration is the equilibrium state at $T_{\rm initial}=\infty$. Any sudden change of temperature therefore drives the system out of equilibrium, and its subsequent relaxation consists in progressively acquiring the spatial correlations characteristic of the final temperature $T_{\rm final}$.

As a result, the nature of this relaxation depends strongly on the relative positions of the initial and final temperatures with respect to T_c . To probe these dynamical processes, we monitor time-dependent observables such as the magnetization and the autocorrelation function:

$$m(t) = \frac{1}{N} \sum_{i=1}^{N} \langle \sigma_i(t) \rangle , \qquad C(t, t_w) = \frac{1}{N} \sum_{i=1}^{N} \langle \sigma_i(t) \sigma_i(t_w + t) \rangle , \qquad (34)$$

where $\sigma_i(t)$ is the spin at site i and time t, and t_w is the waiting time (or age) elapsed after the quench before starting the measurement. The magnetization density m(t) measures the instantaneous global alignment of spins, while the autocorrelation $C(t,t_w)$ quantifies how much of the configuration present at time t_w survives after an additional time t. In simulations, the averages $\langle \cdots \rangle$ are taken over different initial configurations and/or independent thermal histories for statistical accuracy.

To capture spatial structure, we also consider the equal-time spatial correlation function:

$$C(r,t) = \frac{1}{N} \sum_{j=1}^{N} \sum_{k=1}^{N} \langle \sigma_j(t) \sigma_k(t) \rangle \Big|_{|\mathbf{r}_j - \mathbf{r}_k| = r},$$
(35)

which measures the degree of alignment between spins separated by a distance r at a given time t. A characteristic growing length R(t,T) can be extracted from the decay of C(r,t) and

its evolution is central to distinguishing different dynamical regimes after a quench. This is normally defined as

$$R(t) = \frac{\int d^d r \, r^2 \, C(r, t)}{\int d^d r \, r \, C(r, t)} \,, \tag{36}$$

or simply as the distance r at which the correlation function has decay over some characteristic value (i.e. 1/e its value at t=0).

This growing length directly measures the typical size of magnetic domains that are aligned in the same orientation at time t.

Quench from $T_{\rm initial} = \infty$ to $T_{\rm final} > T_c$

In this situation, the random initial configuration and the target temperature share the same type of correlations, so the relaxation is fast. More importantly, the relaxation time τ_R remains finite even in the thermodynamic limit, where the system size tends to infinity. As a result, the autocorrelation decays exponentially,

$$C(t, t_w) \sim e^{-(t - t_w)/\tau_R};$$
 (37)

After this short transient, the magnetization relaxes to its equilibrium value (zero in this case), and the associated correlation length converges to the equilibrium value at the final temperature, $R(t) \to \xi_{\rm eq}(T_{\rm final})$.

The same reasoning applies if the initial configuration is perfectly ordered ($T_{\rm initial} = 0$) and we quench to $T_{\rm final} < T_c$. In this case, the initial correlation length is already effectively infinite, so the quench merely adds thermal fluctuations. The relaxation time is again finite, and the system reaches equilibrium exponentially fast.

Quench from $T_{\rm initial} = \infty$ to $T_{\rm final} < T_c$

This is referred to as a subcritical quench, and its out-of-equilibrium dynamics is well described by the coarsening or domain growth picture.

The initial configuration is completely uncorrelated with the target equilibrium state, and no external magnetic field is present to favor a uniform spin alignment. Instead, spins tend to align locally with their neighbors.

Due to spontaneous symmetry breaking, different regions of the system independently choose one of the two symmetry-related equilibrium states: one with positive magnetization and the other with negative magnetization. This produces a mosaic of magnetic domains of opposite orientation, which subsequently compete and evolve in time. Snapshots of a time evolution of a spin system displaying this domain growth is shown in Fig. 6

The typical domain size grows algebraically:

$$R(t) \sim t^{1/z_d};,\tag{38}$$

with $z_d = 2$ for single-spin-flip dynamics.

In a finite system, thermalization occurs once one orientation takes over the entire sample of volume L^d , i.e., when $R(t) \approx L$. In the thermodynamic limit, $L \to \infty$, domain growth continues without bound and the relaxation time becomes effectively infinite.

Autocorrelations reveal aging—the property that the system's relaxation slows down as it gets older. The decay of autocorrelations depends on both the time difference $t-t_w$ and the waiting time t_w , because $R(t_w)$ at the start of the measurement sets the scale for further evolution. An example of this behavior is shown in Fig. 7.

(i) For short time differences, $t-t_w \ll t_w$, the autocorrelation decays to a quasi-equilibrium value determined by correlations within the domains: $C_{\rm st}(t-t_w) \simeq q_{\rm EA} = m^2$.

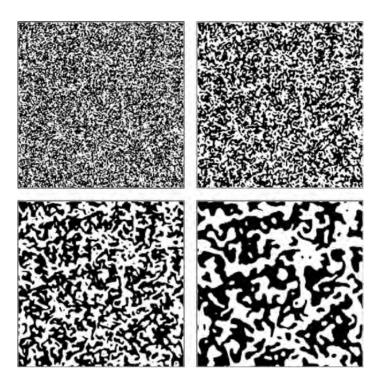


Fig. 6: Snapshots of domain growth in the two-dimensional Ising model with spin-flip dynamics and Metropolis acceptance rates. These snapshots are obtained from a simulation of a $L \times L$ square lattice with periodic boundary conditions and L = 5000. We start from a random initial configuration and perform the quench to $T \simeq 0.73\,T_c$. Image taken from Ref. [48].

(ii) For long time separations, $t-t_w\gg t_w$, the system exhibits aging and the autocorrelation follows the scaling form

$$C_{\rm ag}(t, t_w) \sim f\left(\frac{R(t)}{R(t_w)}\right),$$
 (39)

reflecting the slow restructuring of the domain mosaic.

A compact way to express this stationary + aging decomposition is

$$C(t, t_w) = C_{\rm st}(t - t_w) + C_{\rm ag}\left(\frac{R(t)}{R(t_w)}\right). \tag{40}$$

Quench from $T_{\text{initial}} = \infty$ to T_c

The target equilibrium state is at the continuous phase transition point, where the correlation length is infinite and correlations decay algebraically in space. The system is quenched from a completely disordered configuration (short-range correlations) into a state that is scale-free at equilibrium.

After the quench, the system develops correlated regions whose typical size $R_c(t)$ grows algebraically:

$$R_c(t) \sim t^{1/z_c} \,, \tag{41}$$

with z_c the dynamic critical exponent, that in general depends on the dimensionality of the problem ($z_c \simeq 2.17$ for the 2d Ising model with single-spin-flip dynamics). It is expected to saturate to the value $z_c = z_d = 2$ for $d > D_u = 4$ —with D_u being the upper critical dimension

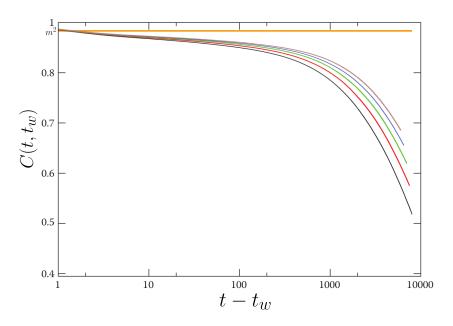


Fig. 7: Plot of $C(t,t_w)$ for a random initial configuration of the 2d Ising model, suddenly quenched to $T_{\rm final} = 0.66\,T_c$. The waiting time increases in steps of 500 from 2000 up to 4000 (bottom to top). The golden horizontal line indicates the EA order parameter $m^2 = 0.97$, corresponding with the equilibrium value within the largest domain. Figure adapted from Ref. [48].

of the Ising model—according to renormalization group calculations [49]. Unlike in the subcritical case, these regions are critical clusters rather than pure up/down domains—their internal correlations already exhibit power-law decay.

Again, in a finite system, equilibrium is reached when $R_c(t) \approx L$, with a relaxation time scaling as

$$au_R \sim L^{z_c};$$
 (42)

In the thermodynamic limit, $L \to \infty$, τ_R diverges and the system never fully equilibrates at finite times. This divergence of τ_R at criticality is the manifestation of *critical slowing down*: as the correlation length grows, it takes increasingly longer for fluctuations to relax, making the dynamics extremely slow near T_c .

Two-time autocorrelations again reveal aging, but here the stationary and aging parts combine multiplicatively,

$$C(t, t_w) = C_{\rm st}(t - t_w) C_{\rm ag}\left(\frac{R_c(t)}{R_c(t_w)}\right),\tag{43}$$

where $C_{\rm st}(\tau) \sim \tau^{-2\beta/(\nu z_c)}$ with β , ν being the Ising critical exponents.

(i) For short time differences, $t-t_w \ll t_w$, the stationary part dominates and the autocorrelation decays as

$$C(t, t_w) \simeq (t - t_w)^{-2\beta/(\nu z_c)},\tag{44}$$

reflecting the quasi-equilibrium decay of critical correlations.

(ii) For long time separations, $t - t_w \gg t_w$, both the stationary and aging parts matter, and one finds

$$C(t, t_w) \simeq t_w^{-b} f_c \left(\frac{R_c(t)}{R_c(t_w)} \right), \qquad b = \frac{2\beta}{\nu z_c},$$
 (45)

In a critical quench to T_c , the growth of $R_c(t)$ under local dynamics leads to large relaxation times, $\tau_R \sim L^{z_c}$, making it computationally expensive to obtain well-equilibrated critical configurations. This phenomenon is known as *critical slowing down*. The Fortuin–Kasteleyn–Coniglio–Klein (FK–CK) cluster construction [50–52] provides a way to overcome this bottleneck.

In the FK–CK framework, spins with the same orientation are grouped into clusters by adding bonds between them with probability

$$p_B = 1 - e^{-2\beta J} \,. \tag{46}$$

This probabilistic construction exactly encodes the equilibrium spatial critical correlations of the ferromagnetic Ising model.

Monte Carlo algorithms such as Swendsen–Wang and Wolff [53–55] exploit this property by using FK–CK clusters as collective update units: entire clusters are flipped in a single stochastic move. This dynamics reproduces the correct critical correlations in a single update step, effectively bypassing the slow coarsening process of local updates and eliminating critical slowing down for unfrustrated systems.

However, the FK–CK construction is not applicable in general to systems with frustration. Even when a ferromagnetic phase exists, the presence of antiferromagnetic bonds can make p_B negative, rendering the cluster-building rule unfeasible for numerical applications. In such cases, no equivalent cluster-based algorithm exists for the generic paramagnetic–ferromagnetic transition in frustrated systems.

Core Idea: FK-CK clusters with frustration

Although negative probabilities arising in frustrated spin systems prevent the direct use of these clusters in numerical simulations, their analytical continuation to negative values enables the FK–CK definition to be explored in exactly solvable models.

The frustrated random-bond Ising model—where the fraction of ferromagnetic and antiferromagnetic bonds is controlled by the parameter ρ —is exactly solvable on certain geometries. In this setting, closed-form analytical expressions can be derived to investigate the properties of these clusters and to determine whether critical correlations persist as ρ is tuned away from the pure ferromagnetic Ising limit.

Such analysis makes it possible to assess whether the original definition continues to capture the critical correlations of the paramagnetic–ferromagnetic transition, even though these clusters remain unusable for direct numerical simulations.

Slow Dynamics in Glassy Systems

Like the paramagnetic–ferromagnetic transition, approaching the spin glass critical temperature $T_{\rm SG}$ from above is accompanied by *critical slowing down*: the dynamical correlation length grows only algebraically in time, leading to long relaxation times. More strikingly, this sluggish dynamics does not disappear in the low-temperature spin glass phase.

In contrast to the ferromagnet, where ordering proceeds rapidly once inside the ordered phase, spin glasses must still build up a complex frozen pattern of correlations whose structure is not known a priori. This makes equilibration extremely slow, and even more subtle phenomena such as *temperature chaos* may come into play: the equilibrium spin configuration changes dramatically under an arbitrarily small temperature variation once the system size exceeds a characteristic overlap length [46, 56–58]. As a result, a quench from $T_1 < T_{\rm SG}$ to $T_2 < T_1$ still exhibits aging and slow relaxation.

As in the ferromagnetic case, the two-time autocorrelation function is the standard probe of aging, now averaged also over disorder realizations:

$$C(t, t_w) = \frac{1}{N} \sum_{i=1}^{N} \left[\langle \sigma_i(t) \sigma_i(t + t_w) \rangle \right], \tag{47}$$

where $[\cdots]$ denotes the average over different quenched disorder samples.

The key difference lies in the spatial correlations. In a spin glass, the emergent order is not captured by the spin–spin correlation function in real space, but rather by correlations between *replicas*. In a given sample, the spins' orientations appear completely random, and no coherent pattern is visible spatially. However, when one considers the *overlaps* between two replicas—configurations with the same disorder realization but independent thermal histories—a clear pattern emerges in the low-temperature phase. This is illustrated in Fig. 8, where the two spin configurations within the spin glass phase look completely random, yet their overlap reveals well-defined correlated regions.

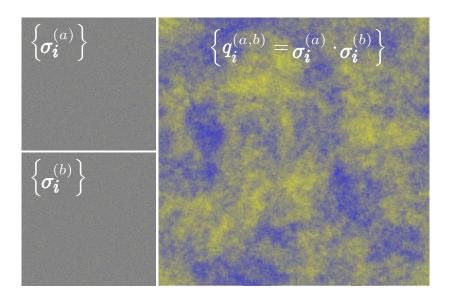


Fig. 8: Illustration of the spin overlap in a spin glass. The left panels show configurations of two independent real replicas of the same disorder sample, evolved for a long time in the low-temperature phase using Monte Carlo dynamics. The color map represents the magnetization in the XY plane, averaged along the Z axis, which vanishes in both replicas. In contrast, the right panel shows the spin overlap defined, where well-defined correlated regions become visible. Adapted from Ref. [59].

This observation motivates the definition of the four-point correlation function.

$$C_4(r,t) = \frac{1}{N} \sum_{\substack{i,j=1\\ |\vec{r}_i - \vec{r}_j| = r}}^{N} \left[\left\langle \sigma_j^{(a)}(t) \sigma_j^{(b)}(t) \, \sigma_i^{(a)}(t) \sigma_i^{(b)}(t) \right\rangle \right] , \tag{48}$$

or, more compactly, as the correlation of the local overlaps

$$C_4(r,t) = \frac{1}{N} \sum_{\substack{i,j=1\\ |\vec{r_i} - \vec{r_j}| = r}}^{N} \left[\langle q_i(t)q_j(t) \rangle \right], \tag{49}$$

with $q_i(t) = \sigma_i^{(a)}(t) \, \sigma_i^{(b)}(t)$. The characteristic length over which $C_4(r,t)$ decays defines the spin glass growing length, estimated as

$$R(t) = 2 \int_0^\infty dr \ C_4(r,t) \ ,$$
 (50)

following the definition of Ref. [60].

A qualitatively similar phenomenon occurs in particle systems that undergo a glass transition where the dynamics become arrested. These systems are known as glass-forming liquids or structural glasses.

Glasses are amorphous solids formed when a liquid is cooled—or compressed—rapidly enough to avoid crystallization. In this state, particles remain disordered as in a liquid, but their dynamics is so slow that the material behaves mechanically like a solid.

The transformation from a supercooled liquid to a glass is called the glass transition T_g . It is not a sharp thermodynamic phase transition, but rather but an empirical quantity: it marks the temperature below which the material becomes so viscous that it effectively stops flowing. This slowdown occurs without any significant change in the spatial structure of the liquid [19].

Upon supercooling, the structural relaxation time τ_R increases extremely rapidly, often by 14–15 orders of magnitude over a temperature range of only about 30% [19]. This dramatic growth is commonly described by the Vogel–Fulcher–Tammann (VFT) law,

$$\tau_R(T) \simeq \tau_0 \exp\left[\frac{A}{T - T_0}\right] ,$$
(51)

where T_0 is the putative divergence temperature at which dynamics would completely freeze if the VFT form held all the way down.

The extreme slowdown of structural relaxation near the glass transition is a major challenge for both experiments and simulations. Standard molecular dynamics or Monte Carlo methods become impractical when τ_R exceeds accessible simulation timescales.

The recently proposed SWAP algorithm [61–63] offered a new way to bypass this bottle-neck. It was originally devised in poly-disperse mixtures of particles—particles of different diameters—that exhibit the same properties of standard glass-formers. It incorporates non-local exchange moves between particles in addition to local displacements of particles generated by Monte Carlo or molecular dynamics moves.

In systems where the method is efficient, SWAP moves may be rarely accepted; however, when they are, the relaxation dynamics is significantly accelerated. In some cases, the SWAP algorithm can achieve speed-ups of up to 10^{10} compared to traditional methods, enabling the sampling of equilibrated configurations at very low temperatures that were previously inaccessible.

In some theoretical scenarios, notably the Random First-Order Transition (RFOT) theory, T_0 is identified with the *Kauzmann temperature* T_K , marking a genuine thermodynamic glass transition where an underlying static order would emerge. RFOT attempts to explain the phenomenology observed in glass-forming liquids through finite-dimensional corrections of mean-field (infinite dimensional) models of spin glasses [64–68].

In these mean-field models, the sluggish dynamics are directly associated to the rugged free-energy landscape presented in Fig. 5 with barriers that become infinite in the thermodynamic limit. RFOT suggests that barriers between states remain finite if the model is considered on a finite-dimensional lattice and that the dramatic slowdown in structural glasses is tied to the existence of a plethora of metastable amorphous states in the free-energy landscape [69]. In this view, the dramatic slowdown of dynamics stems from the growth of static correlations associated with this emerging order. Whether such a transition exists in

real, finite-dimensional systems remains a central open question in glassy physics [19, 21, 70].

Other competing theories propose that the glass transition does not stem from any underlying thermodynamic singularity, but instead emerges purely from dynamical constraints that become increasingly restrictive at lower temperatures or higher densities, where particles cluster and their motion becomes arrested [71, 72]. This perspective is often explored through *kinetically constrained models*—minimalist models devoid of any thermodynamic phase behavior—in which the dynamics are governed by simple rules specifying how individual particles can move within the system [73–75].

In this context, the remarkable acceleration achieved by the SWAP algorithm was later interpreted as direct evidence against static, cooperative explanations of the glass transition, such as that proposed by the RFOT theory [76]. The central argument is that artificially altering the dynamics of a system should not substantially modify its static properties. Since SWAP introduces only artificial dynamical moves yet succeeds in dramatically accelerating relaxation, this suggests that the slow structural relaxation time may not originate from an underlying thermodynamically driven growing length, but rather from purely dynamical constraints that SWAP moves effectively bypass.

Along these lines, recent studies have examined the impact of the SWAP algorithm on the dynamics of kinetically constrained models [77]. In these models, SWAP likewise produces a dramatic acceleration of structural relaxation, reinforcing the view that its efficiency does not depend on growing correlations associated with an emerging thermodynamic order.

Core idea: Testing the SWAP method in spin lattice systems

The SWAP method can be adapted to spin lattice systems by assigning each classical Ising spin a positive-definite amplitude (or length), in direct analogy to the polydispersity of particle sizes in glass-forming systems. These "soft spins" enable non-local exchange Monte Carlo moves that preserve detailed balance and guarantee convergence to equilibrium. In disorder-free systems, such as the standard Ising model, this modification is expected to have little effect. However, in spin systems with quenched disorder—such as the Edwards—Anderson model—non-local exchanges can affect the local energetic constraints, increasing the acceptance rate of Monte Carlo moves and potentially accelerating the relaxation of spin degrees of freedom.

Unlike kinetically constrained models and glass formers, the Edwards–Anderson system lacks dynamical constraints: its slow dynamics arises from the competition between ferromagnetic and antiferromagnetic interactions introduced by quenched disorder. Studying the effect of SWAP in this setting therefore offers a complementary perspective to its investigation in kinetically constrained models.

Quantum Ergodicity

In classical systems, a pure state typically refers to a macroscopic configuration—characterized, for example, by the magnetization m in the ferromagnet, or the overlap q in the spin glass phase. In contrast, a pure state in quantum mechanics refers to a complete description of a system's wavefunction, represented by a state vector in Hilbert space.

Any quantum pure state evolves according to the time-dependent Schrödinger equation

$$\hat{\mathcal{H}} |\Psi(t)\rangle = i\hbar \frac{\partial}{\partial t} |\Psi(t)\rangle , \qquad (52)$$

where $|\Psi(t)\rangle$ is a properly normalized ket-vector in the associated Hilbert space of the

Hamiltonian, i.e. $\langle \Psi(t)|\Psi(t)\rangle=1$. Any pure state can be expressed as a superposition of the eigenstates of the Hamiltonian, whose stationary eigenstates satisfy

$$\hat{\mathcal{H}} |\psi_{\alpha}\rangle = E_{\alpha} |\psi_{\alpha}\rangle . \tag{53}$$

The time evolution of a generic stationary quantum state $|\Psi(0)\rangle$ is given by ⁴

$$|\Psi(t)\rangle = e^{-i\hat{\mathcal{H}}t} |\Psi(0)\rangle = \sum_{\alpha} C_{\alpha} e^{-iE_{\alpha}t} |\psi_{\alpha}\rangle$$
 (54)

Accordingly, the probability of finding the system in a given eigenstate $|\psi_{\alpha}\rangle$ is $p_{\alpha}=|C_{\alpha}|^2=|\langle\psi_{\alpha}|\Psi(0)\rangle$. This probability remains constant over time, as the evolution operator just introduces a phase factor $e^{iE_{\alpha}t}$ that disappears under complex conjugation.

For a local observable $\hat{\mathcal{O}}$ in the Hamiltonian eigenbasis, its expectation value is

$$\langle \hat{\mathcal{O}} \rangle_{t} = \langle \Psi(t) | \hat{\mathcal{O}} | \Psi(t) \rangle = \sum_{\alpha,\beta} C_{\alpha}^{*} C_{\beta} e^{i(E_{\alpha} - E_{\beta})t} \mathcal{O}_{\alpha\beta}$$

$$= \sum_{\alpha} |C_{\alpha}|^{2} \mathcal{O}_{\alpha\alpha} + \sum_{\alpha \neq \beta} C_{\alpha}^{*} C_{\beta} e^{i(E_{\alpha} - E_{\beta})t} \mathcal{O}_{\alpha\beta} ,$$
(55)

where $\mathcal{O}_{\alpha\beta} = \langle \psi_{\alpha} | \mathcal{O} | \psi_{\beta} \rangle$.

We say that $\hat{\mathcal{O}}$ reaches equilibrium if its expectation value evolves from its initial value toward a unique, time-independent value with only small fluctuations over experimental timescales. However, the form of Eq. (55) suggests that a generic many-body pure state cannot equilibrate in general, for several reasons.

- (i) The time-independent diagonal sum depends explicitly on the coefficients C_{α} , which are fixed by the initial state $|\Psi(0)\rangle$
- (ii) The off-diagonal term is oscillatory and, in general, does not decay to zero unless explicitly time-averaged, so the instantaneous expectation value continues to fluctuate.
- (ii) Energy degeneracies generate persistent time-independent off-diagonal contributions, again determined by the initial state.

Historically, these observations prompted the belief that isolated quantum systems cannot thermalize, and that thermalization requires weak coupling to an external bath, making the system effectively open. In such cases, only macroscopic observables matter, and their equilibrium values can be predicted from statistical ensembles without knowledge of the microscopic pure state. This is the standard formulation of quantum statistical mechanics.

In order to describe the equilibrium behavior, we therefore consider averages over ensembles, i.e. \mathcal{M} copies of identical systems. Each member of this ensemble is described by the same Hamiltonian $\hat{\mathcal{H}}$ but starts from a different initial state $|\Psi_k(0)\rangle$, with $k=1,\ldots,\mathcal{M}$. Let p_k denote the fraction of systems in the ensemble that share the same initial condition $|\Psi_k(0)\rangle$, with $\sum_{k=1}^{\mathcal{M}} p_k = 1$. The ensemble average of an observable is then

$$\langle \hat{\mathcal{O}} \rangle_t = \sum_{k=1}^{\mathcal{M}} p_k \langle \Psi_k(t) | \hat{\mathcal{O}} | \Psi_k(t) \rangle$$
 (56)

⁴We will adhere to natural units $\hbar=1$ from this point onward

This in turn can be rewritten in terms of the eigenbasis of Eq. (53)

$$\langle \hat{\mathcal{O}} \rangle_t = \sum_k \sum_\alpha p_k \langle \Psi_k(t) | \hat{\mathcal{O}} | \psi_\alpha \rangle \langle \psi_\alpha | \Psi_k(t) \rangle$$
 (57)

$$= \sum_{\alpha} \langle \psi_{\alpha} | \left(\sum_{k} p_{k} | \Psi_{k}(t) \rangle \langle \Psi_{k}(t) | \right) \hat{\mathcal{O}} | \psi_{\alpha} \rangle \tag{58}$$

$$= \sum_{\alpha} \langle \psi_{\alpha} | \hat{\rho}(t) \hat{\mathcal{O}} | \psi_{\alpha} \rangle = \operatorname{Tr} \left[\hat{\varrho}(t) \hat{\mathcal{O}} \right] = \operatorname{Tr} \left[\hat{\mathcal{O}} \hat{\varrho}(t) \right], \tag{59}$$

where the quantity

$$\hat{\varrho}(t) \equiv \sum_{k} p_k |\Psi_k(t)\rangle \langle \Psi_k(t)| \tag{60}$$

is the density matrix (or operator). It generalizes the description of quantum states to include statistical ensembles, incorporating both unitary time evolution and probabilistic mixing.

By averaging over initial conditions, the explicit dependence on a single set of coefficients C_{α} in Eq. (55) is removed, while the full time evolution is encoded in $\hat{\varrho}(t)$. This resolves the three issues that can prevent equilibration in Eq. (55): the diagonal part no longer retains arbitrary memory of one initial state, oscillatory off-diagonal terms average out in the ensemble, and degeneracy-induced constants are treated naturally within the trace formalism.

Since all time dependence of the average $\langle \hat{\mathcal{O}} \rangle_t$ is contained in $\hat{\varrho}(t)$, we can write down an equation of motion for the density matrix. It follows directly from the Schrödinger equation that

$$i\frac{\partial\hat{\varrho}}{\partial t} = [\hat{\mathcal{H}}, \hat{\varrho}(t)] , \qquad (61)$$

where $[\cdot,\cdot]$ denotes the commutator, and we recall that we are working in natural units $\hbar=1$. This is the quantum-mechanical analogue of Liouville's equation (3).

The density matrix is a Hermitian, positive-semidefinite operator with unit trace, and therefore obeys the following properties:

- (i) $0 \le \hat{\varrho}_{\alpha\alpha}(t) \le 1$, where $\hat{\varrho}_{\alpha\alpha}(t) = \langle \psi_{\alpha} | \hat{\varrho}(t) | \psi_{\alpha} \rangle$ in the eigenbasis of $\hat{\mathcal{H}}$.
- (ii) Tr $\hat{\varrho}(t) = \sum_{\alpha} \hat{\varrho}_{\alpha\alpha}(t) = 1$.
- (iii) $\operatorname{Tr}\left[\hat{\varrho}(t)^2\right] \leq 1$, with equality holding for pure states.

If the physical system we are describing is in thermal equilibrium, the ensemble average must be stationary:

$$\frac{\partial \hat{\varrho}}{\partial t} = 0 \ . \tag{62}$$

From Liouville's equation this implies that $\hat{\varrho}$ must commute with the Hamiltonian, so the two share a common eigenbasis. For a time-independent Hamiltonian, we can therefore write $\hat{\varrho}$ as a function of $\hat{\mathcal{H}}$,

$$\hat{\varrho} = \varrho(\hat{\mathcal{H}}) \ . \tag{63}$$

Expanding in the Hamiltonian's eigenbasis $\{|\psi_{\alpha}\rangle\}$

$$\varrho(\hat{\mathcal{H}}) = \sum_{\alpha} \varrho(E_{\alpha}) |\psi_{\alpha}\rangle \langle \psi_{\alpha}| , \qquad (64)$$

As a result, equilibrium will be described by a diagonal density matrix:

$$\varrho_{\alpha\beta} = \varrho(E_{\alpha})\,\delta_{\alpha\beta} \ . \tag{65}$$

The diagonal element $\varrho_{\alpha\alpha}$ gives the probability of finding the system in eigenstate $|\psi_{\alpha}\rangle$; its functional dependence on E_{α} is fixed by the choice of statistical ensemble. For example, in the microcanonical ensemble all states within a narrow energy shell $[U-\Delta,U]$ are equally probable:

$$\varrho(E_{\alpha}) = \begin{cases} 1/\Omega(U, \Delta) & \text{if } E_{\alpha} \in [U - \Delta, U], \\ 0 & \text{otherwise,} \end{cases}$$
 (66)

where $\Omega(U, \Delta)$ is the number of states in the energy shell.

Thus, even though the microscopic wavefunction of an isolated system evolves unitarily and retains full memory of the initial state, coarse-grained measurements probe a mixed density matrix—typically microcanonical or canonical—whose predictions agree with long-time averages for thermalizing systems. Thermalization, in this sense, was therefore not viewed as a property of individual energy eigenstates, but as a consequence of statistical typicality: almost all states within a narrow energy window are 'thermal' for the purposes of observable physics.

This picture changed dramatically with breakthrough contributions in the 1990s by Deutsch [78] and Srednicki [79–81], who showed that an isolated, non-integrable quantum system can equilibrate when measuring local observables.⁵ The key insight is that, while the entire isolated system retains full information about its initial state, sufficiently small subsystems can act as if they are in thermal equilibrium with the rest of the system. This idea is encapsulated in the *eigenstate thermalization hypothesis* (ETH) [82, 83].

Returning to the time evolution of a closed system in Eq. (55), we now ask under what conditions thermalization can emerge without invoking the ensemble picture. We consider a macroscopically large quantum system with $\mathcal N$ degrees of freedom, and assume that the energy eigenstates with significant weight in the initial pure state $|\Psi(0)\rangle$ lie within a narrow energy window $U-\Delta \leq E_{\alpha} \leq U$. We will focus on the behavior of local observables—such as the number density at a given site i, \hat{n}_i —whose expectation values depend only on a small subset of the system's degrees of freedom.

In the microcanonical ensemble, the equilibrium value of the observable $\hat{\mathcal{O}}$ is

$$\langle \hat{\mathcal{O}} \rangle_{\text{eq}} = \frac{\text{Tr} \left[\hat{\mathcal{O}} \, \theta(U\hat{\mathbb{I}} - \hat{\mathcal{H}}) \right]}{\Omega(U, \Delta)} \,,$$
 (67)

where $\theta(U\hat{\mathbb{I}} - \hat{\mathcal{H}})$ is the Heaviside step function projecting onto the energy shell $[U - \Delta, U]$, and $\Omega(U, \Delta) = \operatorname{Tr} \theta(U\hat{\mathbb{I}} - \hat{\mathcal{H}})$ is the number of states in that shell.

For an observable described by Eq. (55) to thermalize, its long-time value should satisfy

$$\langle \hat{\mathcal{O}} \rangle_t \approx \langle \hat{\mathcal{O}} \rangle_{\mathrm{eq}}$$
 .

This is ensured if the following conditions hold:

(a) The diagonal term matches the equilibrium value,

$$\langle \hat{\mathcal{O}} \rangle_{\text{eq}} = \sum_{\alpha} |C_{\alpha}|^2 \, \mathcal{O}_{\alpha\alpha} \ .$$
 (68)

- (b) The Hamiltonian has no large number of degeneracies, which could otherwise alter (a).
- (c) The off-diagonal term in Eq. (55) decays and fluctuates around zero at long times.

⁵We will comment in more detail on how integrable quantum systems can resist thermalization in Chapter 3 of this thesis.

Remarkably, generic non-integrable many-body Hamiltonians $\hat{\mathcal{H}}$ and most local operators $\hat{\mathcal{O}}$ satisfy these conditions. They are summarized by the *ETH ansatz* for the matrix elements of $\hat{\mathcal{O}}$ in the energy eigenbasis:

$$\mathcal{O}_{\alpha\beta} = \mathcal{O}(\overline{E}) \,\delta_{\alpha\beta} + e^{-S(\overline{E})/2} \,f_{\mathcal{O}}(\overline{E}, \omega) \,R_{\alpha\beta} \,, \tag{69}$$

where:

- $\overline{E}=(E_{\alpha}+E_{\beta})/2$ is the mean energy,
- $\omega = E_{\alpha} E_{\beta}$ is the eigenenergy difference,
- $S(\overline{E}) \sim \ln \Omega(\overline{E})$ is the thermodynamic entropy at \overline{E} ,
- $\mathcal{O}(\overline{E})$ is a smooth function equal to the microcanonical expectation value at energy \overline{E} ,
- $f_{\mathcal{O}}(\overline{E},\omega)$ is a smooth function controlling the size of off-diagonal terms,
- $R_{\alpha\beta}$ is a random variable with zero mean and unit variance.

The ETH implies that, for non-integrable systems, individual many-body eigenstates are already 'thermal' for local observables: $\mathcal{O}_{\alpha\alpha}$ matches the microcanonical prediction at energy E_{α} , while $\mathcal{O}_{\alpha\beta}$ with $\alpha \neq \beta$ vanishes exponentially with system size. In this picture, temporal fluctuations become negligible in the thermodynamic limit, and long-time expectation values for a generic pure state are indistinguishable from those predicted by the microcanoical ensemble.

A central insight of ETH is that thermalization does not require ensemble averaging over many states—the structure of each typical eigenstate is already such that any sufficiently small subsystem behaves as if it were in contact with a thermal bath formed by the rest of the system. This naturally raises the question: how can a closed quantum system, evolving unitarily and remaining in a pure state, exhibit thermal properties in its parts?

The von Neumann entropy, which quantifies the amount of quantum uncertainty in a state, is defined as

$$S[\hat{\varrho}] = -\text{Tr}\left[\hat{\varrho}\ln\hat{\varrho}\right] = -\sum_{\alpha} \lambda_{\alpha} \ln \lambda_{\alpha} , \qquad (70)$$

where λ_{α} are the eigenvalues of $\hat{\varrho}$, satisfying $0 \leq \lambda_{\alpha} \leq 1$. For a pure state of the whole system, $\hat{\varrho} = |\Psi\rangle \langle \Psi|$, the eigenvalue spectrum is $1,0,0,\ldots$, so $S[\hat{\varrho}] = 0$. Since the evolution is unitary, $S[\hat{\varrho}(t)] = S[\hat{\varrho}(0)] = 0$ at all times—meaning the total quantum information of the initial state is exactly preserved.

This demonstrates that an isolated system described by a single pure state will never thermalize as a whole. However, if we partition the degrees of freedom into two disjoint subsystems A and B, the reduced density matrices are defined by partial tracing:

$$\hat{\varrho}_A = \operatorname{Tr}_B \hat{\varrho}, \qquad \hat{\varrho}_B = \operatorname{Tr}_A \hat{\varrho};.$$
 (71)

Although the entropy of the global state remains zero, the reduced states $\hat{\varrho}_A$ and $\hat{\varrho}_B$ are generally mixed when the two subsystems are entangled, and their von Neumann entropies can be nonzero. In systems obeying ETH, the von Neumann entropy of a small subsystem A grows in time and eventually saturates to a value

$$S[\hat{\varrho}_A^{\text{eq}}] = -\text{Tr}\left[\hat{\varrho}_A^{\text{eq}} \ln \hat{\varrho}_A^{\text{eq}}\right],\tag{72}$$

where $\hat{\varrho}_A^{\text{eq}} = \operatorname{Tr}_B \hat{\varrho}_{\text{eq}}$ is obtained by tracing out the degrees of freedom of B from the thermal equilibrium density matrix of the full system. This equilibrium value coincides with the

thermodynamic entropy predicted by statistical ensembles, providing a direct link between ETH and the emergence of thermal behavior in subsystems of isolated quantum systems.

As with the classical ergodic hypothesis, ETH has been rigorously proven only for certain random matrix models and a handful of exactly solvable cases [84]. Nonetheless, the current body of experimental and numerical evidence suggests that ETH behavior manifests for a vast number of macroscopic, non-integrable Hamiltonians [83], yet it remains an open question whether ETH is also a necessary condition for thermalization [83]. As a result, a lot of research has focused to search and understand exceptional cases in which isolated non-integrable systems exhibit the breakdown of ETH.

In particular, Kurchan, Foini, and Pappalardi [85–87] have formulated a generalized ETH that describes not only the thermal bulk of the spectrum but also the statistical properties of rare nonthermal eigenstates. Their work connects ETH to large-deviation theory, showing that fluctuations around the ETH prediction are exponentially suppressed with system size, but can still play an important role in small or mesoscopic systems. As a result, current research often focuses on identifying and understanding such exceptional cases where isolated non-integrable systems exhibit partial or complete breakdown of ETH.

One of the most striking examples of ETH breakdown arises in quantum systems with quenched disorder. In such systems, thermalization can fail due to purely quantum interference effects that prevent particles from exploring the entire system, and thus achieve thermalization. The paradigmatic example is the Anderson insulator [88]—a non-interacting system in which the eigenstates are *localized*, meaning that the probability of finding the particle at a site i, given it was initially at site j, is strongly peaked near i=j and decays exponentially with distance:

$$|\psi(i)|^2 \propto e^{-|i-j|/\xi_{\text{loc}}} \,, \tag{73}$$

where ξ is a characteristic localization length. In this case, a pure state has significant support only on a finite set of sites in real space, and thus only on a small fraction of the Hamiltonian's eigenstates. A key open question is whether this lack of thermalization remains robust in the presence of interactions, potentially giving rise to a stable *many-body localized* (MBL) phase in which ETH is violated.

The possibility of localization in interacting quantum systems was first identified through perturbative analyses by Basko, Aleiner, and Altshuler [89] and by Gornyi, Mirlin, and Polyakov [90]. These works suggested that a closed, interacting many-body system could retain memory of its initial conditions indefinitely, forming a stable dynamical phase—now called *many-body localization* (MBL)—in which the eigenstate thermalization hypothesis (ETH) fails.

Unlike ETH eigenstates, MBL eigenstates are spatially localized and retain a long-term 'memory' of the detailed quantum information encoded in the initial pure state [91–94]. The energy spectrum in the MBL phase contains many gaps, and the matrix elements of local operators deviate strongly from the ETH form (69): the diagonal elements are not smooth functions of energy, while most off-diagonal elements are nearly zero, except for a small fraction that remain large.

In recent years, the stability of the MBL phase has become a subject of intense debate [95–103]. Numerical evidence supporting MBL has been obtained only for relatively small system sizes, due to the exponential growth of the Hilbert space. This makes the results highly sensitive to finite-size effects, manifested for instance in a systematic drift of the apparent transition point, raising the possibility that observed MBL signatures reflect a finite-size crossover rather than a stable phase of matter.

A further challenge comes from theoretical scenarios in which non-perturbative effects—specifically, rare regions of atypically weak disorder—restore ergodic behavior and recover the ETH in the thermodynamic limit [94, 95, 102]. Such rare events challenge the perturba-

tive picture of MBL and suggest that the phase may ultimately be unstable. Determining whether MBL can genuinely persist or is inevitably destroyed in the thermodynamic limit remains a central open problem, yet its resolution is hampered by the severe numerical constraints inherent to many-body quantum systems.

Core idea: Observables capturing the effect of rare events

Large deviation theory and extreme value statistics are standard tools for analyzing classical disordered systems. In some mean-field models of disordered systems, there exists a putative thermodynamic transition driven by rare disorder realizations that generate anomalously large Boltzmann weights. This behavior can be characterized through the *quenched* and *annealed* free-energy densities, defined as

$$f_q(\beta) = \frac{1}{\beta N} \left[\ln \mathcal{Z}(\beta) \right], \qquad f_a(\beta) = \frac{1}{\beta N} \ln \left[\mathcal{Z}(\beta) \right],$$
 (74)

where β is the inverse temperature, N the system size, and $[\cdots]$ denotes the average over disorder realizations.

These two quantities display markedly different behaviors, but coincide in the thermodynamic limit within the high-temperature (replica-symmetric) phase. Because of this, they serve as sensitive observables for estimating the critical transition at finite sizes. An analogous strategy can be applied to the study of the MBL transition, by constructing equivalent observables capable of capturing the influence of rare disorder realizations. Employing such measures may yield valuable insights into the stability of the MBL phase at larger sizes, even when numerical studies are limited to small systems.

The aim of this thesis

This thesis is dedicated to the study of problems that arise in both numerical and theoretical investigations of physical systems with quenched disorder undergoing an ergodicity-breaking transition. We analyze three distinct physical systems, each exhibiting a different class of ergodicity-breaking transition.

As discussed throughout this introduction, the presence of ergodicity-breaking mechanisms hinders the validity of standard theoretical and numerical tools in equilibrium statistical mechanics. When ergodicity breaking mechanisms are present—whether arising from the proximity to a critical point or across an entire phase—the assumptions underlying equilibrium descriptions no longer hold.

For this reason, a significant portion of contemporary statistical mechanics research focuses on bridging the gaps left by equilibrium theory and on developing novel analytical, numerical, and experimental approaches capable of describing scenarios in which ergodicity is broken.

These ergodicity-breaking mechanisms arise in a wide variety of physical systems, particularly within the broad field of disordered systems. In each case, the specific physical context has motivated the development of tailored methods and tools to address the unique challenges it presents.

Nevertheless, the phenomenology associated with ergodicity breaking—such as slow relaxation and the persistence of initial conditions—shows remarkable similarities across otherwise distinct systems. This opens a fertile ground for cross-fertilization, where techniques originally devised in one domain can be adapted to another. Such exchanges not only deepen our understanding of the systems themselves but also test and refine the generality

of the methods employed. In this light, the main aim of this thesis is precisely profit from this cross-fertilization and adapt novel methods—well established in the study of other physical systems—to the specific problems addressed here.

The Physical Systems of Interest

In the Ising model, the phenomenon of critical slowing down near the critical point makes standard simulations with local dynamics extremely slow. However, the spatial critical correlations can be accurately described by the Fortuin–Kasteleyn–Coniglio–Klein (FK–CK) clusters, which provide a theoretical mapping to another physical system (mathematical percolation). This mapping forms the basis for highly efficient numerical algorithms that eliminate critical slowing down at the paramagnetic–ferromagnetic transition.

However, these algorithms—and the exact percolation mapping—do not extend straightforwardly to systems with frustration. In such cases, the construction relies on a probability that becomes negative when both ferromagnetic and antiferromagnetic bonds are present, rendering numerical simulations unfeasible. Nonetheless, for certain models admitting closed-form analytical solutions, one can extend the definition of FK–CK clusters by analytically continuing this probability to negative values. This allows us to test the validity of these clusters in such cases.

We will explore the definition of these clusters in the frustrated random bond Ising model, defined by the Hamiltonian

$$\mathcal{H} = -\sum_{\langle ij \rangle} J_{ij} \sigma_i \sigma_j - h_{\text{ext}} \sum_i \sigma_i , \qquad (75)$$

where σ_i are Ising degrees of freedom, $h_{\rm ext}$ is an external magnetic field and J_{ij} are random bonds uniformly distributed according to the probability distribution

$$\mathcal{P}(J_{ij}) = \rho \delta(J_{ij} - J_0) + (1 - \rho)\delta(J_{ij} + J_0) , \qquad (76)$$

with $J_0 > 0$. We will study this model on the Bethe lattice geometry, that allows for an exact analytical solution. We will introduce and explain in detail the properties of this lattice geometry in Chapter 1.

For $\rho=1$ the ferromagnetic Ising model is recovered, where the critical clusters are described by the FK–CK formalism with a positive bond probability defined by

$$p_B = 1 - e^{-2\beta J_0} \ . {(77)}$$

We investigate how the percolation of FK–CK clusters in this model changes as the fraction of antiferromagnetic bonds increases (ρ < 1), and whether these clusters preserve the correct critical properties of the transition despite involving negative probabilities.

We will also briefly explore alternative cluster definitions proposed to generalize the percolation construction to the paramagnetic–spin-glass transition, evaluating them and pointing out subtle caveats that arise in this mapping and possible interesting future directions to explore.

As mentioned above, even with algorithms that mitigate critical slowing down, the spin-glass phase poses additional challenges: its entire low-temperature regime exhibits extremely sluggish dynamics. A similar situation occurs in structural glasses, where dynamics become arrested at very low temperatures, making it difficult to reach consensus on the nature of the glassy phase. This has motivated the development of novel algorithms based on nonphysical dynamics that nevertheless sample equilibrium configurations efficiently.

Among these, the SWAP algorithm has achieved remarkable speedups—up to 10^{10} in time scales—in particle-based models of glass-forming liquids.

The SWAP algorithm has not only been successful in practical terms but has also become central to debates about the nature of the glass transition, with some arguing that the transition is purely dynamical and that SWAP results support this view.

In this work, we adapt the SWAP algorithm to a finite-dimensional spin-glass system. Specifically, the two-dimensional Edwards–Anderson model whose Hamiltonian reads

$$\mathcal{H}_J(\sigma_i) = -\sum_{\langle ij\rangle} J_{ij}\sigma_i\sigma_j , \qquad (78)$$

and with bimodal distribution of the bonds J_{ij}

$$\mathcal{P}(J_{ij}) = \frac{1}{2}\delta(J_{ij} - J_0) + \frac{1}{2}\delta(J_{ij} + J_0) . \tag{79}$$

Although this model has no finite-temperature spin-glass transition, its zero-temperature ground states exhibit spin-glass order, and relaxation toward them becomes sluggish even within the paramagnetic phase. This makes it an ideal testing ground for adapting SWAP in a spin-lattice context. Our goals are to assess whether the adapted algorithm accelerates dynamics and samples spin-glass ground states effectively, and to clarify the role of dynamical effects in the original particle-based SWAP algorithm.

In order to do this we will study a slight modification, where we include a random positive-definite degree of freedom τ_i that are associated to the spins by $s_i = \sigma_i \tau_i$. This new degree of freedom will play a role analogous to the diameter in the original particle system.

This modification leads to the Hamiltonian

$$\mathcal{H} = -\sum_{\langle ij\rangle} J_{ij}\sigma_i \tau_i \sigma_j \tau_j , \qquad (80)$$

and we have chosen τ_i to be independent and identically distributed random variables, initially drawn from a normalized box distribution, $p_{\tau}(\tau_i)$, *i.e.*

$$\tau_i \in [1 - \Delta/2, 1 + \Delta/2], \qquad 0 \le \Delta \le 2.$$
 (81)

The parameter Δ controls the spin length variation and we will consider $\Delta \leq 2$ to ensure that $\tau_i \geq 0$.

This is the main model of interest and we will refer to it as the Δ -model. The standard Edwards-Anderson model with Ising spins is recovered for $\Delta = 0$.

In the quantum setting, the main challenge is the exponentially large Hilbert space, which limits numerical simulations to very small system sizes. Furthermore, ergodicity-breaking mechanisms often appear to be dominated by rare events that can restore ergodicity, casting doubt on the existence of a true localization transition.

By applying large-deviation methods—commonly used in mean-field classical disordered systems exhibiting a freezing transition—we evaluate the impact of these rare events, enabling conclusions that are more representative of larger system sizes despite working with small ones.

Specifically, we will study the XXZ quantum spin- $\frac{1}{2}$ chain, that serves as the paradigmatic model of the MBL transition. The Hamiltonian, for a chain of L sites, is given by

$$\hat{\mathcal{H}} = \sum_{i=1}^{L} \left(\hat{S}_{i}^{x} \hat{S}_{i+1}^{x} + \hat{S}_{i}^{y} \hat{S}_{i+1}^{y} + \Delta \hat{S}_{i}^{z} \hat{S}_{i+1}^{z} + h_{i} \hat{S}_{i}^{z} \right) , \tag{82}$$

with Δ a constant controlling the strength of the interactions and h_i a random magnetic field uniformly drawn from [-W, W].

In mean-field classical disordered systems, a freezing transition occurs at very low temperatures, where the behavior of the *quenched* and *annealed* free-energy densities differs significantly through the effect of rare events. These free-energy densities are defined as

$$f_q(\beta) = \frac{1}{\beta N} \left[\ln \mathcal{Z}(\beta) \right], \qquad f_a(\beta) = \frac{1}{\beta N} \ln \left[\mathcal{Z}(\beta) \right],$$
 (83)

where β is the inverse temperature, N is the system size, and $[\cdots]$ denotes the average over disorder realizations.

However, both $f_q(\beta)$ and $f_a(\beta)$ coincide in the thermodynamic limit within the high-temperature phase. Consequently, they serve as sensitive observables for estimating the critical transition at finite sizes.

Inspired by this approach, we define analogous quantities for the quantum many-body problem that capture the rare events responsible for restoring ergodicity within the MBL phase. By analyzing the finite-size behavior of these observables as L increases, we estimate the critical disorder strength $W_{\rm MBL}(L)$ at which the ergodicity-breaking transition occurs, explicitly incorporating the effects of these rare events into its determination.

All these studies share a unifying theme: the use of diverse methods within the broader field of disordered systems to tackle problems where standard approaches fail due to the nature of the ergodicity-breaking phase transition. The novel approaches explored here are:

- (i) Percolation mapping to frustrated systems.
- (ii) Non-physical SWAP algorithms for spin-lattice models.
- (iii) Analysis of ergodicity-restoring rare events using large-deviation methods in the many-body localization transition.

Together, these investigations illustrate the value of a cross-fertilization strategy—applying methods across different contexts within disordered systems—to address challenges that transcend specific physical models.

The structure

Each chapter addresses one of three cases, outlining the relevant theoretical framework, the central problem under investigation, and the method used to address it.

Chapter 1: The Critical Clusters of Frustrated Spin Systems. This chapter explores the Fortuin–Kasteleyn–Coniglio–Klein (FK–CK) clusters and related definitions in the context of the frustrated random-bond Ising model on the Bethe lattice. It begins with an intuitive introduction to cluster formation in the Ising model, together with the basics of percolation theory and the Bethe lattice geometry. The original formulation of FK–CK clusters is reviewed, followed by an extension of this framework to analytically continued clusters with negative probabilities, applied to the frustrated random bond Ising model. The main equations are derived, solved, and used to determine the properties of the FK–CK clusters in this context. Finally, alternative definitions of critical clusters proposed for the paramagnetic-to-spin-glass transition are discussed, and one of these is explored on the Bethe lattice using the same approach.

Chapter 2: SWAP Dynamics for Frustrated Spin Systems. This chapter adapts the SWAP Monte Carlo algorithm—originally developed for particle systems—to spin-lattice systems, focusing on the two-dimensional Edwards—Anderson model. After a brief review of structural glass phenomenology and two main competing theoretical frameworks for the glass transition, the original SWAP method is introduced, along with its relevance to this debate. The adaptation of SWAP to spin systems is then presented, including tests on simpler models such as the Ising model to assess the dynamical effects of the method and to confirm that equilibrium properties remain essentially unchanged. The method is subsequently applied to the 2d Edwards—Anderson model, a frustrated spin system with spin-glass ground states.

Chapter 3: The Importance of Rare Events in Many-Body Localization. This chapter investigates rare events that restore ergodicity within the many-body localized phase of interacting, isolated quantum systems, using the XXZ spin chain in a random field as a case study. It begins with a review of Anderson localization and its theoretical framework, followed by an overview of many-body localization and an explanation of it ergodicity breaking mechanism. The XXZ chain is introduced as the central model, together with standard numerical observables used to probe the MBL phase, highlighting their limitations. Non-perturbative mechanisms proposed to restore ergodicity are then discussed and explained, they are referred to as 'thermal avalanches' and 'many-body resonances'. The novel observable along with the main large deviation method is presented to better capture the role of rare events in determining the transition, and its predictions for the XXZ chain are reported. The chapter concludes with results characterizing these events and an interpretation of the MBL transition related to the freezing transition present in some mean-field classical disordered systems.

Chapter 1

The critical clusters of frustrated spin systems

Contents

1.1	The spin configurations across the phase transition	. 34
1.2	Geometrical percolation	
	1.2.1 The Bethe lattice	. 37
	1.2.2 Site percolation on the Bethe lattice	. 40
1.3	The FK-CK correlated bond percolation	
	1.3.1 The FK approach	
	1.3.2 The CK approach	
1.4	The FK-CK clusters for frustrated models	
	1.4.1 The model	
	1.4.2 The FK approach to the frustrated RBIM	
1.5		
	1.5.1 Cavity field and recursion equations	
	1.5.2 Phase diagram	
1.6	Percolation of the FK–CK clusters	
	1.6.1 Ising model	
	1.6.2 Frustated RBIM	
1.7	Spin glass clusters	
	1.7.1 Two-replica FK–CK clusters	
	1.7.2 Multiple-replica FK-CK clusters	
1.8	Conclusions and outlook	

Introduction

The Ising model occupies a central position in statistical physics as the paradigmatic model for studying phase transitions and critical phenomena [104]. It serves as the standard example for understanding the rich and universal features that emerge near continuous (second-order) phase transitions, and nucleation phenomena in first-order phase transitions in the presence of an external magnetic field. Beyond its original formulation, the Ising model has proven instrumental in elucidating the link between thermodynamic phase transitions and the geometrical organization of microscopic constituents [105, 106].

The phase transition is marked by the emergence of spin clusters with well-defined geometric and statistical characteristics, which encapsulate the essential information about spatial correlations among the spin degrees of freedom at criticality. A rigorous formulation of such clusters was first established in the early 1970s by Kees Fortuin and Piet Kasteleyn [50, 51], who introduced the random cluster model—an exact reformulation of the Potts (and, in particular, Ising) model partition function in terms of weighted sums over bond-connected clusters. This framework not only unified the description of percolation and Potts models but also offered a probabilistic representation of the Ising model's critical behavior through correlated cluster formation. This approach has since had a profound influence on stochastic geometry and mathematical physics [107].

About a decade later, Antonio Coniglio and William Klein (CK) independently demonstrated that the critical clusters of the Ising model could be constructed by introducing a bond probability that connects parallel spins [52, 108]. This probabilistic construction, was later generalized to other models by Robert Edwards and Alan Sokal [109].

Collectively, these developments established a coherent and formal correspondence between percolation phenomena and thermodynamic phase transitions, through the identification of critical clusters—referred to as FK–CK clusters—that percolate exactly at the thermodynamic critical temperature T_c .

Moreover, the geometrical critical exponents—such as those associated with cluster size and cluster-size distributions—are directly related to the well-known thermodynamic critical exponents of the Ising model. This means that the shapes and size distributions of clusters directly reflect the system's thermodynamic transition. The definition of FK–CK clusters was later extended [110] to the antiferromagnetic Ising model [111] and to Ising models with ferromagnetic interactions beyond nearest neighbors [112].

This geometrical insight has not only deepened our theoretical understanding of phase transitions but has also inspired the development of highly efficient cluster-based algorithms for Monte Carlo simulations, such as the Swendsen-Wang [54, 113] and Wolff algorithms [55]. These methods effectively suppress the phenomenon of critical slowing down by constructing non-local updates based on the FK–CK clusters, thereby drastically improving simulation performance near criticality [114, 115].

Nevertheless, one key challenge of this formulation is that is not easily translated into frustrated systems of competing interactions. A naive implementation of the FK–CK is unfeasible in numerical simulations, as the introduction of frustration necessarily leads to negative probabilities. This has led to extensive research aimed at finding suitable alternatives to the FK–CK clusters that can reproduce the critical behavior of these systems. In spin glass systems, for example, new proposals have been put forward [116–121] as alternative candidates to replace the FK–CK definition. None of these, however, reproduce the critical behavior for arbitrary dimensionality nor produce algorithms as efficient as the ones developed for the unfrustrated cases [122, 123].

In this chapter, we will revisit the original FK–CK definition of clusters and use it to study frustrated systems. Specifically, for the frustrated random bond Ising model on the

Bethe lattice. Although in these systems the FK–CK definition becomes intractable in Monte Carlo simulations—due to the emergence of negative probabilities—one can still probe the percolation of these clusters in exactly solvable models, such as the one considered here. We will first explore whether the original FK–CK clusters can still capture the critical behavior associated with the thermodynamic transition. In addition, we will investigate alternative procedures and definitions, analyzing their limitations as well as potential avenues for further research.

The structure of this chapter is as follows. In Sec. 1.1, we provide an intuitive introduction to the cluster formulation of the Ising model. Once this motivation is established, Sec. 1.2 presents the formal framework of percolation theory, introducing its fundamental concepts and the main geometry used throughout this work—the Bethe lattice. The history of the FK–CK clusters is reviewed in Sec. 1.3, where we distinguish between the formulations originally introduced by Fortuin and Kasteleyn, and by Coniglio and Klein—two approaches that, as we will see, are equivalent. In Sec. 1.4, we extend this framework to frustrated systems, using the frustrated random-bond Ising model as a case study. Then, in Sec. 1.6, we present the solution of FK–CK percolation on the Bethe lattice, along with additional results, including an alternative cluster definition. This alternative definition sets the stage for discussing the current efforts of mapping the percolation transition onto the model when spin glass order is present, which is the main focus of Sec. 1.7. Finally, Sec. 3.8 concludes with a summary of the main results, as well as an outline of possible future extensions.

1.1 The spin configurations across the phase transition

Consider the Ising model in absence of an external magnetic field, with N spin variables σ_i interacting with their nearest neighbors through a constant ferromagnetic bond J>0. The Hamiltonian is given by

$$\mathcal{H} = -J \sum_{\langle i,j \rangle} \sigma_i \sigma_j \ . \tag{1.1}$$

This model exhibits a continuous phase transition, from a disordered paramagnetic phase at high-temperatures to an ordered ferromagnetic phase at low-temperatures. The transition is identified by the gobal magnetization density

$$m = \frac{1}{N} \sum_{i=1}^{N} \langle \sigma_i \rangle , \qquad (1.2)$$

where $\langle \cdots \rangle$ denotes the average over the Gibbs-Boltzmann distribution. As explained in the Introduction Chapter of this thesis, at high-temperatures $(T > T_c)$ the magnetization density vanishes m = 0, while at low temperatures $(T < T_c)$ the Z_2 is spontaneously broken and the system decides between a positively and negatively magnetized equilibrium states.

In Fig. 1.1 we show three examples of an equilibrium 2d spin configuration in the square lattice. The configuration at $T>T_c$ is in the paramagnetic phase (left panel), in which the entropic term of the free-energy dominates, and the configuration seems completely random. For $T=T_c$, the configuration is at criticality, where the magnetic domains acquire a self-similar structure and the system becomes scale-invariant (center panel). Finally, the right-most panel shows a configuration in the ferromagnetic phase $(T< T_c)$, where the system has chosen the positive magnetized equilibrium state—with some small thermal fluctuations of negative spins.

Coming from the low-temperature phase ($T < T_c$), thermal fluctuations (visible as yellow, negative spins) gradually increase in number and form larger magnetic domains as the

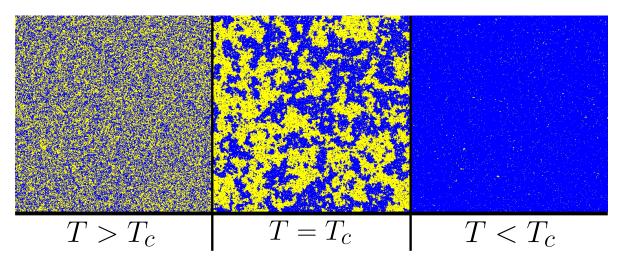


Fig. 1.1: Snapshots of equilibrium configurations for the d=2 Ising model in a square lattice, with positive spins in blue colors and negative spins in yellow.

temperature rises. At the critical point $T=T_c$, the clusters of these minority spins (yellow) have grown to a size comparable to that of the entire system. The same picture can be viewed from the opposite direction: in the high-temperature phase, positive spins (blue) form magnetic domains that grow and eventually span the entire system at some temperature $T \leq T_c$. From these perspectives, one can define a percolation temperature T_p , at which a cluster of geometrical spins —clusters of neighboring spins sharing the same orientation—first spans the entire system. This may correspond either to negative-spin domains emerging as T increases from the ordered phase, or to positive-spin domains forming as T decreases from the disordered, high-temperature phase.

What remains less certain is whether this percolation transition coincides exactly with the ferromagnetic-paramagnetic transition i.e. $T_p = T_c$, and whether the statistical properties of these clusters carried relevant statistical information about the thermal correlations of the microscopic spins. The most straightforward intuition is that clearly, these clusters should be the *geometrical clusters* meaning, clusters formed by neighboring spins that point in the same direction. In this interpretation,

It was later recognized, independently by Fortuin and Kasteleyn [50, 51] and by Coniglio and Klein [52], that geometrical clusters are too large and fail to accurately capture the thermal fluctuations near the critical point. Instead, the correct critical behavior is recovered when clusters are formed following a probabilistic construction that depends on the bond strength J and the system's temperature T. These are known as FK–CK clusters. A brief summary of the basic framework of percolation theory is provided below, followed by a detailed description of the procedure to construct these clusters.

1.2 Geometrical percolation

Before diving deeper into the FK–CK cluster definition, we will first review the two most simple forms of the percolation problem [124, 125]. Our general setup will be a graph (or lattice) G = (V, E) formed by vertices (or sites) $i \in V$ (with |V| = N) and edges (or bonds) $(ij) \in E$ connecting neighboring sites i, j.

Site percolation

The most basic formulation of the problem considers each site of the lattice as either occupied (i.e., $n_i = 1$), with probability p, or empty (i.e., $n_i = 0$), with probability 1 - p. We are

interested in answering the question: what is the critical value of the probability p that produces a system-spanning—or percolating—cluster in the graph? This is termed as *site* percolation. An example of the transition for a d=2 square lattice is shown in Fig. 1.2. Below a critical threshold probability p_c occupied sites are sparse. However, upon increasing the occupation probability over the threshold p_c , occupied sites proliferate and a system-spanning cluster arises for $p \geq p_c$.

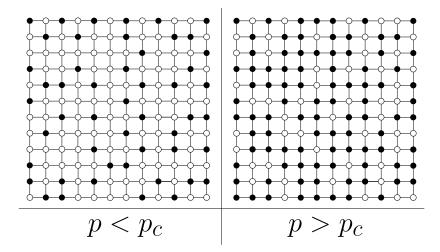


Fig. 1.2: Sample of a 2d square lattice undergoing site percolation.

Bond percolation

The second formulation is concerned with the bonds instead. We consider every site to be occupied $n_i=1 \ \forall i$, but the edges $(ij)\in E$ have a probability p of connecting a pair of sites i,j and a probability 1-p of being absent. Two sample configurations for $p< p_c$ and $p>p_c$ are shown in Fig. 1.3 for the square lattice. In this case, the sites themselves are not explicitly represented, since they are all trivially occupied. Again, at $p\geq p_c$ a percolating cluster arises, identified by the sites connected by bonds extending throughout the entire lattice.

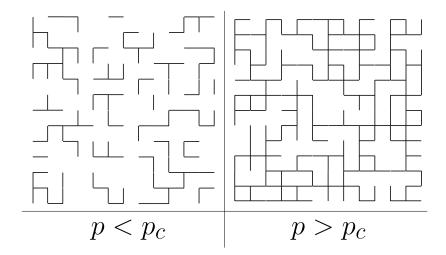


Fig. 1.3: Sample of a 2d square lattice undergoing bond percolation.

As many problems in statistical physics, the critical percolation threshold p_c represents an average property, as the construction of the clusters is itself a random process. For finite

sized systems, the critical probability is the average

$$p_c \equiv \frac{1}{N_S} \sum_{k=1}^{N_S} p_c^{(k)} , \qquad (1.3)$$

with N_S samples—or realizations—of the random process and $p_c^{(k)}$ the respective threshold found for the k-th sample. A key feature of this problem is that sample-to-sample fluctuations diminish with increasing system size, allowing p_c to converge to a well-defined value in the thermodynamic limit $N \to \infty$.

Consequently, one can formally identify a phase transition in which this emergent giant cluster appears as a function of the control parameter p. Several other observables can be used as order parameters detecting the transition. One such observable, is the density of sites belonging to the largest cluster, this is given by

$$n_{\text{max}}(p, N) \equiv \frac{N_{\text{max}}(p, N)}{N} , \qquad (1.4)$$

For $p < p_c$, the size of the largest cluster $(N_{\max}(p,N))$ grows sub-linearly with system size and therefore tends to zero with growing N. In contrast, for $p > p_c$, the scaling becomes extensive $(N_{\max}(p,N) \sim O(N))$ so the normalized size of the largest cluster $(n_{\max}(p,N))$ approaches a finite, non-zero value as N increases.

In the thermodynamic limit $(N \to \infty)$, fluctuations vanish. For $p > p_c$, the largest cluster corresponds exactly to the infinite percolating cluster, and the density of sites within this cluster matches the probability that a randomly chosen site belongs to it. This probability is often referred to as strength or weight of the percolating cluster, here we will denote this quantity as the percolation probability:

$$P = \lim_{N \to \infty} n_{\text{max}}(p, N) , \qquad (1.5)$$

that, as anticipated, serves as an order parameter for the percolation transition:

$$P = 0 \quad \text{for } p \le p_c,$$

$$P \ne 0 \quad \text{for } p > p_c.$$
(1.6)

The percolation threshold p_c depends on the lattice geometry, its dimensionality and on the type of percolation you are considering—either site, bond or a combination of both.

Apart from the order parameter that we have presented in Eq. (1.6), there are several relevant observables for the percolation transition: a correlation length ξ that diverges when the percolation transition is approached, the mean (and largest) cluster size $S(S_{\text{max}})$, the distribution of cluster sizes $S(S_{\text{max}})$, among others. Similarly to second-order phase transitions, these quantities display universal behavior when approaching the critical point $S(S_{\text{max}})$ and universal critical exponents can be defined for them.

There is a very small number of models in which an exact value for p_c and the associated critical exponents can be obtained. Often one needs to rely on numerical simulations, renormalization group calculations, rigorous upper-lower bounds and series of expansions to estimate these quantities. Some of these critical exponents for the site percolation problem are shown in Table 1.1 for regular geometries in d = 1, d = 2 and d = 3.

1.2.1 The Bethe lattice

From the results in Table 1.1, we see that although the case d=1 is exactly solvable, it displays only trivial critical behavior. We therefore turn to the next simplest analytically

Exponent and law		d=2	d=3
p_c	1	0.5927	0.3116
ν Correlation length ($ξ$)			
$\xi \sim p - p_c ^{-\nu}$	1	4/3	0.876(1)
β Percolation probability (P)			
$P \sim p - p_c ^{eta}$	0	5/36	0.4181(8)
γ Mean cluster size (S)			
$S \sim p - p_c ^{-\gamma}$	1	43/18	1.80(5)
σ Largest cluster size (S_{max})			
$S_{ m max} \sim p - p_c ^{-1/\sigma}$	1	36/91	0.445(10)
$ au$ Cluster-size distribution (n_S)			
$n_S \sim S^{-\tau}$	2	187/91	2.189(2)
ω Correction to scaling (ω)			
$\operatorname{std}(\mathcal{O}) \sim N^{y_{\mathcal{O}}}(1 + cN^{-\omega})$	N/A	0.77(4)	0.64(5)

Table 1.1: Critical threshold and exponent values for site percolation in dimensions d = 1, 2, 3. The correction-to-scaling behavior has been established for an arbitrary observable \mathcal{O} with corresponding critical exponent $y_{\mathcal{O}}$. For details see [124, 125].

solvable case within site or bond percolation: the percolation problem set on the Bethe lattice. This is the same geometry in which we will analyze the percolation of FK–CK clusters, and is thus central to the present work.

The Bethe lattice is an infinite, regular tree where every vertex has the same connectivity $\kappa + 1$ —with κ being the branching number of the tree. It has an inherent self-similar structure with no root, no center, and no boundary. This property—along with its tree nature—allows for exact solutions in the form of recursive equations, making a useful playground to understand the properties of many statistical physics models [126, 127].

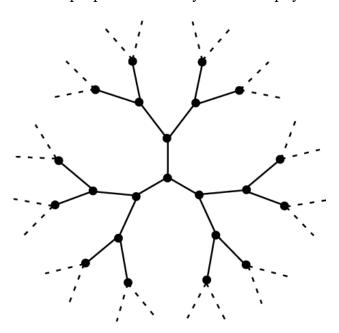


Fig. 1.4: Schematic portion of the infinite Bethe lattice with $\kappa + 1 = 3$. The dashed lines represent other tree segments extending to infinity.

A finite portion of a Bethe lattice with $\kappa + 1 = 3$ is shown in Fig. 1.4. Starting from any given site, the number of accessible sites within n edges on the Bethe lattice grows exponentially with n, in contrast to the polynomial growth $\sim n^d$ observed in a regular

d-dimensional lattice. As a result, the Bethe lattice is often regarded as representing the $d \to \infty$ limit of the problem, effectively yielding a mean-field-like model. Unlike fully connected models, however, it preserves a local structure, which provides a more realistic approximation of spatially extended systems.

This behavior can be related to that of the Cayley tree, a similar tree structure. In this case, each site also has a fixed connectivity of $\kappa+1$, but the lattice possesses a boundary whose sites have connectivity 1. As a result, a central—or root—site can now be properly identified in this structure, as it is shown in Fig. 1.5.

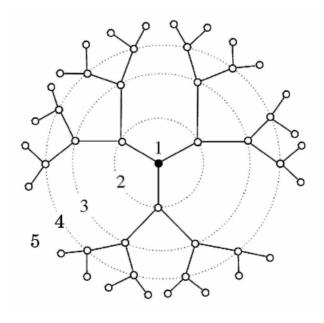


Fig. 1.5: Schematic portion of a Cayley tree with $\kappa + 1 = 3$. The dashed lines represent other tree segments extending to infinity.

The first generation—the central node—branches into $\kappa+1$ sites, which form the second generation. Each of these sites then branches into κ new sites, defining the third generation, and this process continues recursively until the boundary is reached. The total number of nodes in a Cayley tree with n generations therefore is

$$N = 1 + (\kappa + 1) + \kappa(\kappa + 1) + \kappa^{2}(\kappa + 1) + \dots + \kappa^{n-1}(\kappa + 1)$$

$$= \frac{\kappa^{n}(\kappa + 1) - 2}{\kappa - 1},$$
(1.7)

while the surface boundary is composed of $N_{\Omega} = \kappa^{n-1}(\kappa+1)$ sites. In the thermodynamic limit, the fraction of sites in the boundary is

$$\lim_{N \to \infty} \frac{N_{\Omega}}{N} = \frac{\kappa - 1}{\kappa} \,, \tag{1.8}$$

and thus remains finite. Contrary to what is observed in finite-dimensional lattices, where the surface-to-bulk ratio vanishes for large N.

Although the local structure of the Bethe lattice is equivalent to that of the Cayley tree, for models in which boundary effects play an important role, the comparison between both types of lattices should be made with caution [128]. If one wishes to perform numerical simulations whose results can be directly cross-checked with analytical calculations on the Bethe lattice, it is preferable to consider random regular graphs, which share the same fixed connectivity $\kappa+1$ but lack a boundary.

For finite sizes, RRGs more accurately model bulk behavior. An RRG is a graph $G_N = (V, E)$ of size N = |V|, where edges between vertices are placed at random under the constraint that each vertex has a fixed connectivity $\kappa + 1$. These graphs are locally tree-like but contain loops whose typical length grows as $O(\log N)$. Because the Bethe lattice is loopless, the presence of loops in RRGs leads to deviations from Bethe lattice predictions at finite N. These finite-size effects alter observable values, with corrections for local quantities typically scaling as 1/N. Nonetheless, in the infinite-size limit $(N \to \infty)$, RRGs locally converge to the Bethe lattice. This convergence allows critical quantities—such as p_c —to be accurately approximated by their Bethe lattice values when N is large. In contrast, when considering finite-size Cayley trees, the boundary size remains comparable to that of the bulk even in the thermodynamic limit, which can affect system properties and make them less representative of bulk behavior than RRGs.

In this way, the Bethe lattice is not just an abstract mean-field model, but also serves as a benchmark for quantifying finite-size effects in numerical simulations on RRGs, and it qualitatively captures the behavior of more realistic systems.

1.2.2 Site percolation on the Bethe lattice

The general description of site percolation still holds in this case: each site of the lattice is occupied with probability p and empty with probability p. A percolation transition occurs at the critical probability p_c , marked by the emergence of a system-spanning cluster.

Paths are formed by edges connecting occupied site, in this context, having a percolating cluster implies that there is an infinite path formed by occupied sites. If we choose a site on the lattice that is connected to the infinite path by one of its $\kappa+1$ neighbors, there are κ branches leading to a new site, which could be occupied with probability p. Hence, for an arbitrary site connected to a (semi-)infinite path by one of its neighbors, the probability that this path continues through the remaining neighbors is p κ . Consequently, the condition to ensure that there is always an occupied neighboring site—and thus an infinite path—is given by the condition p $\kappa=1$. The critical occupation probability is

$$p_c = \frac{1}{\kappa} \,. \tag{1.9}$$

This argument is also valid where instead of empty sites the access to a neighboring site is blocked by the absence of an edge, making Eq. (1.9) also valid for bond percolation.

We now proceed to define the order parameter, the percolation probability P. This quantity is defined as the probability that any randomly chosen site belongs to the percolating cluster, thanks to the self-similar structure of the Bethe lattice, we can find an exact recursive equation that can be solved to obtain P.

An arbitrary site must necessarily be occupied in order to belong to the percolating cluster. However, this is not a sufficient requirement, as it also needs that (at least) one of its $\kappa+1$ nearest-neighbors is also occupied and connected to the percolating cluster through another one of its subsequent neighbors, continuing this process ad-infinitum.

We begin by considering a specific site in the lattice, and label it by i. This site is occupied with probability p_i , or empty with probability $1-p_i$. Site i has $\kappa+1$ nearest neighbors, collectively denoted by the set ∂i . We have two main probabilities associated to site i: The percolation probability P_i , which represents the probability that site i belongs to the percolating cluster, and the probability that Q_i for site i to not belong to the percolating cluster, here denoted as Q_i .

There are two ways in which the site i may not belong to the percolating cluster: (i) it is not occupied with probability $1 - p_i$ or (ii) it is occupied but it is completely isolated,

meaning that none of its neighbors belong to the percolating cluster, hence:

$$Q_i = 1 - p_i + p_i \prod_{l \in \partial i} Q_{l \to i} , \qquad (1.10)$$

Here, $Q_{l\rightarrow i}$ denotes the probability that site l does not belong to the percolating cluster in the absence of site i. This is one of the simplest examples of the cavity method [35, 127]. The basic idea is that by 'isolating' site i, we can analyze the contribution from its nearest neighbors independently.

It is important to note that $Q_{l\rightarrow i}$ is fundamentally different from Q_i . However, by recursively applying this reasoning to successive generations, we can establish a recursive relation. After neglecting site i, we move forward to one of its other κ neighbors, which we label as j. Site j then has κ available branches leading to new sites, and each of those sites, in turn, connect to κ further sites—excluding j itself. Thus we can construct:

$$Q_{j\to i} = 1 - p_j + p_j \prod_{l\in\partial j\setminus i} Q_{l\to j} . \tag{1.11}$$

In this context, the quantities $Q_{j\to i}$ and $Q_{l\to j}$ are commonly referred to as cavity probabilities, and the omitted site is called the cavity site. Restoring the site independence $p_j = p$, Eq. (1.11) becomes

$$Q = 1 - p + p Q^{\kappa} , (1.12)$$

that can be solved numerically, or exactly for the specific case of $\kappa=2$. We obtain two solutions Q=1 (corresponding to the case $p< p_c$) and Q=(1-p)/p (for $p>p_c$). For an arbitrary site, the probability that it is occupied but not connected to the percolating cluster is p-P that in turn can be expressed in terms of Q as pQ^3 —the probability of the site being occupied but its $\kappa+1=3$ neighbors disconnected from the infinite cluster. Thus $P=p(1-Q^3)$, that together with the previous solution for Q we obtain

$$P = p \left(1 - \frac{(1-p)^3}{p^3} \right) . {(1.13)}$$

Expanding Eq. (1.13) around p_c , one can obtain the critical behavior associated to the exponent β , on the Bethe lattice this corresponds to $\beta = 1$. The same type of calculation can be performed to find several of the observables mentioned previously for the finite-dimensional cases. In Table 1.2, the list of some of them is shown for the Bethe lattice at any connectivity.

These exponents are somewhat analogous to the mean-field exponents of the Ising model, which can be derived using either the molecular field approximation or the Bethe-Peierls approximation. In particular, the fully-connected graph represents a case where the molecular field approximation becomes exact. Similarly, the Bethe lattice corresponds to the geometry in which the Bethe-Peierls approximation is exact. For this reason, we refer to both cases as mean-field-like models.

Any attempt to map the paramagnetic-to-ferromagnetic transition of the Ising model onto a percolation transition should reproduce not just the correct critical temperature T_c , but also the Ising universality class. Consequently, the critical clusters should not belong to the percolation universality class characterized by the exponents in Table 1.1 or Table 1.2, but instead exhibit the corresponding Ising critical exponents, some of these are shown in Table 1.3 for the same geometries and dimensions mentioned above, with the exception of d=1, where it has been replaced to the most interesting case the Bethe lattice, with its mean-field exponents.

Exponent and law		Bethe lattice
p_c		
		$1/\kappa$
ν	Correlation length (ξ)	
	$\xi \sim p - p_c ^{-\nu}$	1
β	Percolation probability (P)	
	$P \sim p - p_c ^{eta}$	1
$\overline{\gamma}$	Mean cluster size (S)	
	$S \sim p - p_c ^{-\gamma}$	1
σ	Largest cluster size (S_{max})	
	$S_{ m max} \sim p-p_c ^{-1/\sigma}$	1/2
$\overline{\tau}$	Cluster-size distribution (n_S)	
	$n_S \sim S^{-\tau}$	5/2

Table 1.2: Critical threshold and exponent values for percolation on the Bethe lattice. For details see [124, 125].

Exponent and law	Beth lattice	d=2	d=3
T_c (in units of J/k_B)	$1/ \operatorname{atanh}(1/\kappa)$	2.269	4.5115(2)
ν Correlation length (ξ)			
$\xi \sim T - T_c ^{-\nu}$	1/2	1	0.6301(4)
β Spontaneous magnetization (m)			
$M \sim T_c - T ^{eta}$	1/2	1/8	0.3264(2)
γ Susceptibility (χ)			
$\chi \sim T - T_c ^{-\gamma}$	1	7/4	1.2372(5)
δ Critical isotherm			_
$m \sim h_{ m ext}^{1/\delta}$	3	15	4.789(2)
α Specific heat (C)			
$C \sim T - T_c ^{-\alpha}$	0	0 (log)	0.110(1)
η Correlation function exponent	0	1/4	0.0363(1)
ω Correction to scaling (ω)	1	2	0.83(3)

Table 1.3: Critical exponents for the Ising universality class: mean-field (Bethe lattice, z=6), 2D square lattice, and 3D cubic lattice. T_c for the Bethe lattice is given by $2J/\ln(z-1)$, yielding $T_c \approx 3.478$ for z=6.

1.3 The FK-CK correlated bond percolation

From the thermodynamic point of view, the main quantities of interest are those which can be obtained directly from the free-energy. However, geometrical structures provide powerful insight of the microscopic mechanisms driving the phase transition, and are also amenable to experimental observations.

Consider the Ising model in its lattice gas form, using the transformation $\sigma_i = 2n_i - 1$, where n_i denotes the occupation of the *i*-th site in the lattice.

$$\mathcal{H} = -\epsilon \sum_{\langle i,j \rangle} n_i n_j \tag{1.14}$$

with $\epsilon > 0$ serving as an attractive interaction between particles, with $\epsilon = 4J$ when compared to Eq. (1.1). Now, the order parameter is given by the average density of particles, that can be written in terms of the magnetization density as (1+m)/2.

In this representation, the system undergoes a low-density to high-density transition as we lower the temperature. This explicit structural change prompted researches to conceive the conglomerates of particles—formed by nearest-neighbor pairs—as the central objects driving the phase transition i.e. those describing the thermal correlations. This is completely equivalent as conceived in the spin case, shown in Fig. 1.1.

This idea was further motivated by numerical simulations in d=2, that showed evidence of the mean cluster size S diverging at the critical temperature T_c [129]. However, this enthusiasm faded away when numerical simulations at d=3 [130] and the exact solution on the Bethe lattice [131] showed that the percolation of these clusters, occurred at a different temperature $T_p \neq T_c$. Specifically, the geometrical clusters appear to be larger than the critical clusters carrying the correct spin-spin (or density-density) correlations.

Instead of forming clusters with all occupied neighboring sites (or parallel spins) we reduce their size by creating new clusters, referred to as Ising droplets. Among all pairs of neighboring parallel spins—related to the ferromagnetic equilibrium state—we add bonds between them with probability

$$p_B = 1 - e^{-\beta \epsilon/2} = 1 - e^{-2\beta J}$$
, (1.15)

yielding in a correlated bond percolation problem.

For a spin configuration at low temperatures, we show a comparison of the percolating phase between Ising droplets and geometrical clusters in Fig. 1.6. This pictorial representation uses the same convention presented in Fig. 1.1, where blue vertices denote positive spins, and yellow ones negative spins serving as thermal fluctuations in the ordered phase. On the left panel, the cluster is formed solely by positive whereas the Ising droplets are a subset of the latter, constructed with a random process controlled by the probability p_B .

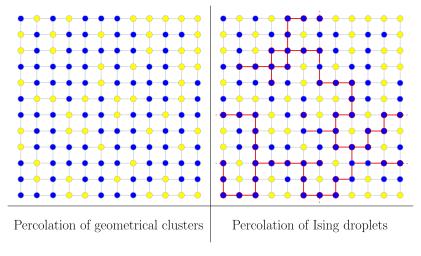


Fig. 1.6: A percolating Ising spin configuration ($T < T_p$) for (left panel) geometrical clusters and (right panel) Ising droplets.

The mean size of the Ising droplets diverges at the Ising critical point with Ising exponents, as was shown by Coniglio and Klein (CK) in their seminal article [52]. It was later shown [109], that these Ising droplets have the same statistics as the clusters introduced by Fortuin and Kasteleyn [50, 51] (FK) in their random cluster model. Within their framework, it was possible to prove that these Ising droplets carry the critical spin-spin (or density-density) correlations, meaning that

$$\langle \sigma_i \sigma_j \rangle = \langle \gamma_{ij}^{\parallel} \rangle. \tag{1.16}$$

where $\langle \sigma_i \sigma_j \rangle$ is the spin-spin correlation function, and $\langle \gamma_{ij}^{\parallel} \rangle$ denotes the probability that spins i and j are parallel and belong to a same (random) cluster of the corresponding percolation problem. Consequently, it follows that

$$|\langle \sigma_i \rangle| = \langle \gamma_i \rangle$$
, (1.17)

with $\langle \gamma_i \rangle$ being the probability of the spin *i* belonging to the percolating cluster.

This geometric identification then enables updating several correlated spins (or particles) all at once, thus accelerating configurational sampling down to the critical temperature, T_c . This was the breakthrough of cluster-based numerical algorithms, proposed by Swendsen and Wang [132], and Wolff [55], that are able to perform collective moves avoiding the critical slowing down seen in standard Monte Carlo simulations.

1.3.1 The FK approach

Consider a system of Ising spins on a lattice with ferromagnetic nearest-neighbor interactions, for which we have set the ground state energy to zero—by adding an irrelevant constant—the Hamiltonian is

$$\mathcal{H}(\{\sigma_i\}) = -\sum_{\langle i,j\rangle} J(\sigma_i \sigma_j - 1). \tag{1.18}$$

We can modify this Hamiltonian by 'diluting' the interactions i.e. replace it by

$$\mathcal{H}'(\{\sigma_i\}) = -\sum_{\langle i,j\rangle} J'_{ij}(\sigma_i \sigma_j - 1), \qquad (1.19)$$

where

$$J'_{ij} = \begin{cases} J' & \text{with probability } p_B, \\ 0 & \text{with probability } 1 - p_B. \end{cases}$$

For a given J', the parameter p_B is chosen such that the Boltzmann weight associated to the original spin configuration, coincides with a spin configuration of the new diluted Hamiltonian, i.e.

$$e^{\beta J(\sigma_i \sigma_j - 1)} = p_B e^{\beta J'(\sigma_i \sigma_j - 1)} + (1 - p_B)$$
 (1.20)

In the limit $\beta J' \to \infty$, we have $e^{\beta J'(\sigma_i \sigma_j - 1)} = \delta_{\sigma_i, \sigma_j}$, and from Eq. (1.20) we get that

$$p_B = 1 - e^{-2\beta J} \,. \tag{1.21}$$

As a result, the Boltzmann factor of each spin configuration is

$$e^{-\beta \mathcal{H}(\{\sigma_i\})} = \prod_{\langle i,j\rangle} e^{\beta J(\sigma_i \sigma_j - 1)} = \prod_{\langle i,j\rangle} \left[p_B \delta_{\sigma_i,\sigma_j} + (1 - p_B) \right] . \tag{1.22}$$

Performing the products in the relation above we can write

$$e^{-\beta \mathcal{H}(\{\sigma_i\})} = \sum_{C} W_{FK}(\{\sigma_i\}, C), \qquad (1.23)$$

where

$$W_{FK}(\{\sigma_i\}, C) = \prod_{\langle i,j\rangle \in C} p_B \delta_{\sigma_i,\sigma_j} \prod_{\langle i,j\rangle \notin C} (1 - p_B)$$

$$= p_B^{|C|} (1 - p_B)^{|A|} \prod_{\langle i,j\rangle \in C} \delta_{\sigma_i,\sigma_j}. \tag{1.24}$$

Here, C is a subset of all the bonds that correspond to a specific configuration of the interactions J', such that the bonds with $J' = \infty$ belong to the cluster configuration C, and the subset of bonds with J' = 0 defines A, with |C| + |A| = |E| for E the set of all bonds. In other words, $W_{FK}(\{\sigma_i\}, C)$ is the statistical weight of a spin configuration $\{\sigma_i\}$ with the set of interactions $\{J'_{ij}\}$ in the diluted model with |C| edges interacting with infinite strength and all other edges interacting with zero strength. The Kronecker delta indicates that two spins connected by an infinite strength interaction must be in the same state. Therefore, the cluster configuration C can be decomposed in clusters of parallel spins connected by infinite strength interactions.

The partition function \mathcal{Z} is then obtained by summing the Boltzmann factor in Eq. (1.23) over all possible spin configurations. Because each disconnected cluster in the cluster configuration C gives a contribution of 2, one then gets

$$\mathcal{Z} = \sum_{C} 2^{N_C} p_B^{|C|} (1 - p_B)^{|A|}, \qquad (1.25)$$

where N_C is the number of clusters in C. Put differently, the FK–CK formalism gives a partition function, $\mathcal{Z} = \sum_C W(C)$, whose structure is equivalent to that of a correlated bond percolation model,

$$W(C) = \sum_{\{\sigma_i\}} W_{FK}(\{\sigma_i\}, C) = 2^{N_C} p_B^{|C|} (1 - p_B)^{|A|},$$
(1.26)

which coincides with the weight of the random bond percolation except for the extra factor 2^{N_C} . All clusters and their weights for the spin configuration $(s_t, \sigma_r, s_b, \sigma_l) = (\uparrow, \uparrow, \uparrow, \uparrow)$ are shown in Fig. 1.7. Clearly, all percolation quantities in this correlated bond model weighted according to equation Eq. (1.26) coincide with the corresponding percolation quantities of the FK clusters made up of parallel spins connected by infinite-strength interaction, whose statistical weight is given by Eq. (1.24).

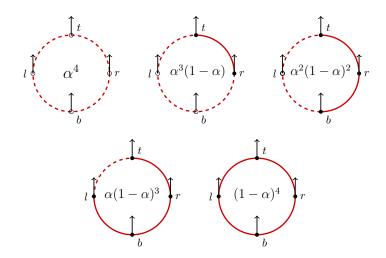


Fig. 1.7: Chain of four Ising spins with ferromagnetic interactions, $J_{ij}=J>0$, under periodic boundary conditions, in its minimal energy spin configuration $(\sigma_t,\sigma_r,\sigma_b,\sigma_l)=(\uparrow,\uparrow,\uparrow,\uparrow)$. All possible clusters are shown, with zero to four links (solid red lines), thus identifying spins as being either part (filled circle) or not (unfilled circle) of a cluster. Cluster multiplicities are 1, 4, 6, 4, and 1, respectively, and FK cluster weights $W_{\rm FK}$ (center) are expressed for bonding probability $\alpha=e^{-2\beta J}\in(0,1]$.

1.3.2 The CK approach

In the approach followed by Coniglio and Klein [52], given a configuration of spins, one introduces fictitious random bonds between neighboring parallel spins with probability p_B , while anti-parallel spins are *not* connected, with probability 1. Clusters are defined as maximal sets of parallel spins connected by bonds. Unlike with the FK approach the interaction energy is left unchanged.

For a given realization of bonds we distinguish the subsets C and A_1 as the sets of neighboring parallel spins respectively connected and not connected by bonds and the subset A_2 of neighboring antiparallel spins. Clearly, $|C| + |A_1| + |A_2| = |E|$ where E is the set of all bonds of the lattice. The statistical weight of a given configuration of spins $\{\sigma_i\}$ and bonds C is given by

$$W_{CK}(\{\sigma_i\}, C) = p_B^{|C|} (1 - p_B)^{|A_1|} e^{-\beta \mathcal{H}(\{\sigma_i\})}.$$
(1.27)

The partition function is obtined by summing Eq (1.27) over all possible spin and bond configurations

$$\mathcal{Z} = \sum_{\{\sigma_i\}} \sum_{C} W_{CK}(\{\sigma_i\}, C) = \sum_{\{\sigma_i\}} e^{-\beta \mathcal{H}(\{\sigma_i\})}$$

$$\tag{1.28}$$

where the second equality follows from $\sum_{C} p_{B}^{|C|} (1 - p_{B})^{|A_{1}|} = 1$.

The partition function of course does not depend on the value of p_B , which controls the bond density. Instead, p_B tunes the size of the clusters formed by these bonds. For example, taking $p_B=1$ the clusters would coincide with nearest-neighbor parallel spins, while for $p_B=0$ the clusters are reduced to single spins.

By choosing $p_B = 1 - e^{-2\beta J}$ (equivalent to the result of Eq. (1.21) in the FK approach) and recognizing that $e^{-\beta \mathcal{H}(\{\sigma_i\})} = e^{-2\beta J|A_2|}$, Eq. (1.27) simplifies to

$$W_{CK}(\{\sigma_i\}, C) = p_B^{|C|} (1 - p_B)^{|A|}, \qquad (1.29)$$

where $A = A_1 \cup A_2 = E \setminus C$. From this equation, we can calculate the weight of the bond configuration C. This bond configuration can occur among all possible spin configurations, however we are interested in taking into account those compatible with the condition of spins being parallel, i.e.

$$W(C) = \sum_{\{\sigma_i\}} W_{CK}(\{\sigma_i\}, C) \prod_{\langle ij \rangle \in C} \delta_{\sigma_i \sigma_j} . \tag{1.30}$$

Inserting the explicit form of Eq.(1.29) we get

$$W(C) = \sum_{\{\sigma_i\}} p_B^{|C|} (1 - p_B)^{|A|} \prod_{\langle ij \rangle \in C} \delta_{\sigma_i \sigma_j} = 2^{N_C} p_B^{|C|} (1 - p_B)^{|A|} , \qquad (1.31)$$

coinciding with the result of the FK approach of (1.26) yielding an equivalent partition function

$$\mathcal{Z} = \sum_{C} 2^{N_C} p_B^{|C|} (1 - p_B)^{|A|} \,. \tag{1.32}$$

The CK clusters and the KF clusters have different meanings. In the CK formalism, the clusters are defined directly in the Ising model as parallel spins connected by fictitious bonds, while in the KF formalism clusters are defined in the equivalent diluted model as parallel spins connected by infinite-strength interactions. However, both approaches are equivalent, as can be seen from the equality between the resulting weights of Eq. (1.26) and Eq. (1.26).

From these statistical weights it follows that, any percolation quantity $\mathcal{O}(C)$ which depends only on the bond configuration C has the same average in both approaches

$$\langle \mathcal{O}(C) \rangle_{FK} = \langle \mathcal{O}(C) \rangle_{CK} = \langle \mathcal{O}(C) \rangle_{W} , \qquad (1.33)$$

where $\langle \cdots \rangle_W$ is the average over bond configurations in the bond-correlated percolation with weights given by W(C). Consequently, it also follows that

$$\langle \sigma_i \sigma_j \rangle = \langle \gamma_{ij}^{\parallel} \rangle_W \,, \tag{1.34}$$

where $\langle \cdots \rangle$ is the thermodynamic average with the Boltzmann weights. Here $\gamma_{ij}^{\parallel}(C) = 1$ if i and j are in the same cluster, and 0 otherwise.

1.4 The FK-CK clusters for frustrated models

In models with frustrated interactions, the bond probabilities p_B can become negative, prompting the search for extended definitions of p_B . We will mainly consider the case in which this frustration manifests by the presence of both ferromagnetic (J > 0) and anti-ferromagnetic bonds (-J).

In a natural generalization for these cases, FK–CK clusters may be constructed by adding bonds between satisfied spin pairs i.e. to those sites i, j with $J_{ij}\sigma_i\sigma_j > 0$. It can be shown that [116, 133]

$$\langle \sigma_i \sigma_j \rangle = \langle \gamma_{ij}^{\parallel} \rangle - \langle \gamma_{ij}^{\parallel} \rangle, \tag{1.35}$$

where $\langle \gamma_{ij}^{\parallel} \rangle$ and $\langle \gamma_{ij}^{\parallel} \rangle$ denote the probability that parallel and antiparallel spins i and j are in the same cluster, respectively. Some of the bonds in the clusters then link neighboring parallel spins and others anti-parallel ones. In this context, generalizing cluster schemes by summing the contribution of parallel and antiparallel spins necessarily overestimate correlations, i.e.,

$$|\langle \sigma_i \sigma_j \rangle| \le \langle \gamma_{ij}^{\parallel} \rangle + \langle \gamma_{ij}^{\parallel} \rangle, \tag{1.36}$$

and the resulting clusters percolate at temperatures higher than T_c .

Therefore, in the critical regime these clusters are largely ineffective at configurational sampling, as has been repeatedly demonstrated in numerical simulation [122, 134–139].

In what follows, we adopt the FK approach to demonstrate that FK–CK clusters still capture the relevant critical correlations in frustated systems; that is, the relation $\langle \sigma_i \sigma_j \rangle = \langle \gamma_{ij}^{\parallel} \rangle_W$ continues to hold. However, for spin pairs connected by an *antiferromagnetic* bond, one must take $p_B < 0$, thereby relinquishing any probabilistic interpretation of this quantity.

We first present the frustrated spin model for which we will provide the proof: the frustrated random bond Ising model (RBIM), also known as the $\pm J$ Ising model.

1.4.1 The model

The main model for this chapter will be the frustrated RBIM, defined by the Hamiltonian

$$\mathcal{H}_{\text{RBIM}}(\{\sigma_i\}) = -\sum_{\langle i,j\rangle} J_{ij}\sigma_i\sigma_j - h_{\text{ext}} \sum_i \sigma_i , \qquad (1.37)$$

where the couplings J_{ij} are taken from a bimodal distribution with probability ρ of being ferromagnetic, $J_{ij} = J_0 > 0$, and probability $1 - \rho$ of being anti-ferromagnetic, $J_{ij} = -J_0$. In other words, the J_{ij} s are taken at random from the probability distribution:

$$\mathcal{P}(J_{ij}) = \rho \delta(J_{ij} - J_0) + (1 - \rho)\delta(J_{ij} + J_0) , \qquad (1.38)$$

with J_0 setting the unit of energy. Clearly, the standard ferromagnetic Ising model is recovered for $\rho=1$, and the standard antiferromagnetic Ising model is recovered for $\rho=0$. There also exists a duality: $\rho \leftrightarrow (1-\rho) \wedge J_0 \leftrightarrow -J_0$.

The RBIM on d=2 [140–147] and d=3 [122] cubic lattices, in particular, has been extensively studied. For all $d\geq 2$ and ρ close to unity the paramagnetic-ferromagnetic transition is Ising-like. As ρ decreases, however, dimensional differences emerge. In particular, in d=3 the model exhibits a finite-temperature spin-glass phase for $\rho\lesssim 0.78$, while in d=2 spin-glass ordering is only present at T=0 (for $\rho\lesssim 0.897$). In mean-field models, such as the infinite-connectivity Sherrington-Kirkpatrick (SK) model as well as comparable models defined on the Bethe lattice, an additional phase emerges: the ferromagnetic-spin-glass (FSG) phase. This intermediate phase between the ferromagnetic and the spin-glass phases exhibits spin-glass behavior while maintaining a non-zero magnetization [148–151].

Interestingly, for all $d \ge 2$, local gauge symmetry gives rise to a Nishimori line [152, 153]

$$\beta_N J_0 = \frac{1}{2} \ln \frac{\rho}{1 - \rho},\tag{1.39}$$

where $\beta_N=1/k_BT_N$ is the inverse Nishimori temperature. Along this line certain thermodynamic quantities, such as the internal energy, have closed-form expressions. This line is also invariant under renormalization-group transformations, as is the paramagnetic-ferromagnetic transition line. The intersection point of the two lines therefore gives rise to a multicritical fixed point; this Nishimori point (NP) separates the paramagnetic, ferromagnetic and spin-glass phases in $d \geq 3$. In d = 2, the dynamical (critical) exponent of single spin flip Metropolis Monte Carlo at the Nishimori point has been estimated to have a very high value, $z \simeq 6$, independently of lattice geometry and random bond distribution [154, 155]. Such a high z significantly inhibits configurational sampling (and equilibration) in this regime.

1.4.2 The FK approach to the frustrated RBIM

Once again, apart from an irrelevant additive constant (and in absence of external magnetic field), the Hamiltonian in Eq. (1.37) can be rewritten as

$$H_{\text{RBIM}}(\{\sigma_i\}) = -\sum_{\langle i,j\rangle} J_{ij}(\sigma_i \sigma_j - 1). \tag{1.40}$$

We then replace this Hamiltonian with an annealed diluted Hamiltonian

$$H'(\lbrace \sigma_i \rbrace) = -\sum_{\langle i,j \rangle} J'_{ij}(\sigma_i \sigma_j - 1), \qquad (1.41)$$

where

$$J_{ij}' = \begin{cases} J' & \text{with probability } p_B^{(ij)} \,, \\ 0 & \text{with probability } 1 - p_B^{(ij)} \,. \end{cases}$$

For a fixed J', $p_B^{(ij)}$ is chosen such that

$$e^{\beta J_{ij}(\sigma_i \sigma_j - 1)} = p_B^{(ij)} e^{\beta J'(\sigma_i \sigma_j - 1)} + \left(1 - p_B^{(ij)}\right)$$
 (1.42)

for each bond. In the limit $J'\to\infty$, we have $e^{\beta J'(\sigma_i\sigma_j-1)}=\delta_{\sigma_i,\sigma_j}$, and $p_B^{(ij)}$ is given by

$$p_B^{(ij)} = 1 - e^{-2\beta J_{ij}}. (1.43)$$

(As anticipated, for $J_{ij} < 0$ the parameter $p_B^{(ij)}$ is negative; the formal construction can nevertheless be continued.) Consequently, the Boltzmann factor is

$$e^{-\beta H_{\text{RBIM}}(\{\sigma_i\})} = \prod_{\langle i,j \rangle} \left[p_B^{(ij)} \delta_{\sigma_i,\sigma_j} + \left(1 - p_B^{(ij)} \right) \right] , \qquad (1.44)$$

and hence we can write

$$e^{-\beta H(\{\sigma_i\})} = \sum_{C} W_{FK}(\{\sigma_i\}, \{J_{ij}\}, C), \qquad (1.45)$$

where

$$W_{FK}(\{\sigma_i\}, \{J_{ij}\}, C) = \prod_{\langle i,j \rangle \in C} p_B^{(ij)} \delta_{\sigma_i, \sigma_j} \prod_{\langle i,j \rangle \notin C} (1 - p_B^{(ij)}). \tag{1.46}$$

Here, C is a subset of all the bonds that correspond to a specific configuration of the interactions J'_{ij} , such that the bonds with $J' = \infty$ belong to the cluster configuration C, and the subset of bonds with J' = 0 defines A, with |C| + |A| = |E| for E the set of all bonds. Note that we have here included the dependence on the full realization of the random couplings $\{J_{ij}\}$ in the definition of the FK statistical weights, because the quenched disorder over bond types make their 'probabilities' differ for each system realization.

The partition function \mathcal{Z} can then be obtained by summing over all spin configurations,

$$\mathcal{Z} = \sum_{C} 2^{N_C} \prod_{\langle i,j \rangle \in C} p_B^{(ij)} \prod_{\langle i,j \rangle \notin C} (1 - p_B^{(ij)}), \qquad (1.47)$$

where N_C is the number of clusters in C. Therefore, the FK–CK formalism gives a partition function, $\mathcal{Z} = \sum_C W(C)$, whose structure is equivalent—albeit, as mentioned above, with some negative bond probabilities—to that of a correlated bond percolation model

$$W(C) = \sum_{\{\sigma_i\}} W_{FK}(\{\sigma_i\}, \{J_{ij}\}, C)$$

$$= 2^{N_C} \prod_{\langle i,j \rangle \in C} p_B^{(ij)} \prod_{\langle i,j \rangle \notin C} (1 - p_B^{(ij)}).$$
(1.48)

It follows that

$$\langle \sigma_i \sigma_j \rangle = \langle \gamma_{ij}^{\parallel} \rangle_W \,, \tag{1.49}$$

where $\gamma_{ij}^{\parallel}(C)=1$ if i and j are in the same cluster, and 0 otherwise. We emphasize that this equality holds when averaging over Boltzmann weights and averaging over bond configurations, for *any* fixed disorder realization, $\{J_{ij}\}$. Consequently, the equality must also hold after averaging over the quenched disorder. A concrete example is discussed in the concluding Sec. 3.8.

Note that the equality between the spin–spin correlation function and $\langle \gamma_{ij}^{\parallel} \rangle_W$ follows from the FK–CK clusters ensuring that the Boltzmann weight of every spin configuration coincides with the statistical weight of the corresponding random-bond percolation model defined by the measure W(C). This property crucially implies that, even if one constructs another cluster model for which clusters percolate exactly at the Ising critical point and exhibit the same critical exponents—as for the α -parameter clusters that will be discussed below—the equality between the Boltzmann weight of the spin configurations and the statistical weight of the cluster model no longer holds. As a result, throughout the phase diagram the spin–spin correlation function differs from the percolation correlation function. Consequently, an algorithm based on such a cluster model is not expected to weaken the critical slowing down with any significance.

1.5 The RBIM on the Bethe lattice

The presence of negative probabilities makes the definition of FK–CK clusters unsuitable for Monte Carlo simulation studies. However, we can forgo the probabilistic interpretation of the quantities p_B and instead study the clusters in models that admit closed analytical expressions, such as the model defined on the Bethe lattice.

In this section, we provide the solution of the frustrated RBIM on the Bethe lattice, along with its thermodynamic phase diagram. Before proceeding to the percolation analysis of the FK–CK clusters, we first examine their behavior in the standard Ising model, corresponding to the case $\rho=1$. This serves as a consistency check for our methods, ensuring that we recover the known results. We then solve the percolation equations for the frustated RBIM and demonstrate that the FK–CK clusters indeed percolate at the critical point, exhibiting the corresponding Ising critical exponents.

1.5.1 Cavity field and recursion equations

As anticipated in Sec. 1.2.2, we will use the cavity method to solve this model on the Bethe lattice. In this case we are not dealing with percolation probabilities, but with marginal probabilities of the microscopic spin configuration. In what follows, we will denote positive and negative spins with \uparrow and \downarrow , respectively.

We will label a current site with o, and label its $\kappa + 1 = 3$ nearest neighbors as i, j and k, as illustrated on the left-hand side of Fig. 1.8. The marginal probability of a microscopic spin configuration at the site o can be parametrized in terms of an effective local field h_o^{eff} as

$$\eta_o(\sigma_o) = \frac{e^{\beta h_o^{\text{eff}} \sigma_o}}{2 \cosh(\beta h_o^{\text{eff}})}.$$
(1.50)

This effective field encapsulates the influence of the neighboring spins on site o. Alternatively, this marginal probability can be expressed explicitly in terms of the interactions with neighboring sites

$$\eta_o(\sigma_o) = \frac{e^{\beta h_{\text{ext}}\sigma_o}}{\mathcal{Z}_{\text{site}}} \prod_{l \in \partial_o} \sum_{\sigma_l} e^{\beta J_{ol}\sigma_o\sigma_l} \eta_{l \to o}(\sigma_l) , \qquad (1.51)$$

where $\mathcal{Z}_{\text{site}}$ is a normalization constant ensuring $\eta_o(\uparrow) + \eta_o(\downarrow) = 1$. The quantity $\eta_{l \to o}(\sigma_l)$ is the cavity marginal configuration probability for the site $l \in \{i, j, k\}$, computed in absence of site o. This cavity construction is depicted on the right-hand side of Fig. 1.8, where the contribution from the site i depends on the remaining neighbors $\{m, n\} = \partial i \setminus o$. This cavity configuration probability satisfies a recursive equation of the same form

$$\eta_{l \to r}(\sigma_l) = \frac{e^{\beta h_{\text{ext}} \sigma_l}}{\mathcal{Z}_{\text{cav}}} \prod_{p \in \partial l \setminus r} \sum_{\sigma_p} e^{\beta J_{lp} \sigma_l \sigma_p} \eta_{p \to l}(\sigma_p) , \qquad (1.52)$$

for any generic site l, with r denoting the cavity site. The factor \mathcal{Z}_{cav} ensures normalization, $\eta_{l\to r}(\uparrow)+\eta_{l\to r}(\downarrow)=1$. Analogously to Eq. (1.50), the cavity configuration probability can also be parametrized using a cavity field $h_{l\to r}$

$$\eta_{l\to r}(\sigma_l) = \frac{e^{\beta h_{l\to r}\sigma_l}}{2\cosh(\beta h_{l\to r})}.$$
(1.53)

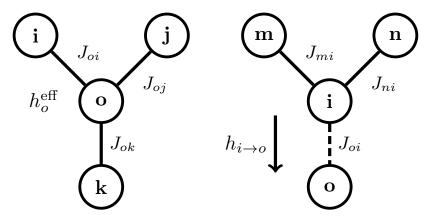


Fig. 1.8: Labeling scheme for a generic frustrated RBIM configuration on a Bethe lattice with connectivity $\kappa+1=3$. (Left) The site o is subjected to an effective field $h_o^{\rm eff}$, arising from interactions with its three nearest neighbors i,j, and k, via couplings $J_{oi},J_{oj},$ and $J_{ok},$ respectively. (Right) The contribution of site i to $h_o^{\rm eff}$, denoted $h_{i\to o}$, reflects the influence of its own neighboring sites m and n, computed in the absence of site o.

With both effective and cavity configuration probabilities expressed in this form, we can derive equivalent expressions for their corresponding fields. The effective field at site o is given by:

$$h_o^{\text{eff}} = h_{\text{ext}} + \frac{1}{\beta} \sum_{l \in \partial o} \operatorname{atanh}(\tanh \beta J_{ol} \tanh \beta h_{l \to o}),$$
 (1.54)

while the cavity field $h_{l\to r}$, which captures the influence on site l in the absence of site r, satisfies the recursion

$$h_{l\to r} = h_{\text{ext}} + \frac{1}{\beta} \sum_{p \in \partial l \setminus r} \operatorname{atanh}(\tanh \beta J_{lp} \tanh \beta h_{p\to l}). \tag{1.55}$$

Note that both effective and cavity fields are themselves random variables, due to the disorder present in the couplings $\{J_{0l}\}$ and $\{J_{lp}\}$ whenever $\rho \neq 1$. The stationary distribution for the cavity fields in Eq. (1.55) can be obtained with arbitrary precision using the population dynamics algorithm [127]. Numerically, this scheme involves initializing a random sample—or population—of size \mathcal{M} of the cavity fields $\{h_{l\rightarrow r}\}$ and iteratively updating each member of this population according to Eq. (1.55). Each member is updated \mathcal{N} times until the entire population converges to a stationary probability distribution. Once the distribution of the cavity fields is known, the effective fields defined in Eq. (1.54) can be computed. Knowing the distribution of effective fields, $\{h_o^{\rm eff}\}$, enables the calculation of all relevant thermodynamic quantities. In particular, the average magnetization density is given by

$$m \equiv [\langle \sigma_o \rangle] = [\tanh(\beta h_o^{\text{eff}})],$$
 (1.56)

where $\langle \cdots \rangle$ denotes the thermal average with respect to the Gibbs–Boltzmann distribution, and $[\cdots]$ represents the average over disorder realizations. In this context, averaging over disorder is equivalent to averaging over the distribution of effective fields $\{h_o^{\rm eff}\}$.

1.5.2 Phase diagram

Interestingly, for h=0, closed-form analytical expressions for the critical temperatures of the paramagnetic-to-ferromagnetic (P-F) and paramagnetic-to-spin-glass (P-SG) transitions

can be obtained [153, 156]

$$\frac{J_0}{T_c} = \operatorname{atanh}\left(\frac{1}{\kappa(2\rho - 1)}\right),$$

$$\frac{J_0}{T_{SG}} = \operatorname{atanh}\left(\frac{1}{\sqrt{\kappa}}\right),$$
(1.57)

respectively. These expressions are calculated from the linear stability of the first two moments of the cavity fields in Eq. (1.55) around zero. In other words, they correspond to the temperatures for which

$$[h_{l\to r}] = 0$$
 and $[h_{l\to r}^2] = 0,$ (1.58)

respectively. The critical dilution at which these critical temperatures coincide is

$$\rho_{\star} = \frac{1}{2} \left(1 + \frac{1}{\sqrt{\kappa}} \right) , \qquad (1.59)$$

which corresponds to a multicritical point that separates a regime in which the system undergoes a paramagnetic-to-ferromagnetic transition (ρ close to 1) from another one in which the ferromagnetic phase is replaced by a phase with spin-glass ordering (ρ close to 1/2). The three phases meet at this point. Coming from T large, the P-F transition corresponds to the temperature at which the system develops a non-zero magnetization, given by Eq. (1.56). By contrast, during the P-SG transition, the spontaneous magnetization remains zero, but the system exhibits a non-zero Edwards–Anderson (EA) order parameter,

$$q_{EA} \equiv [\langle \sigma_o \rangle^2] = [\tanh^2(\beta h_o^{\text{eff}})]. \tag{1.60}$$

A phase with both non-zero magnetization and EA order parameter can be identified, denoted here as the ferromagnetic spin glass (FSG). The transition lines between SG-FSG and FSG-F can be estimated numerically, at the replica-symmetric level, using population dynamics. Coming from the spin glass phase, the critical line SG-FSG is estimated at the onset of finite magnetization, while the FSG-F transition line is determined from the stability of the EA order parameter, measured as

$$\delta q_{EA} = |q_{EA} - q_{ab}|,\tag{1.61}$$

where q_{ab} is the overlap of two-replicas a and b, i.e.

$$q_{ab} = \left[\langle \sigma_o^{(a)} \rangle \langle \sigma_o^{(b)} \rangle \right] = \left[m^{(a)} m^{(b)} \right], \tag{1.62}$$

and $\delta q_{EA}=0$ in the purely ferromagnetic phase. Numerically, this corresponds to two populations of the cavity fields being updated simultaneously, with different initial conditions. A crucial point is that both populations a and b are evolved concomitantly in the population dynamics algorithm [157]. Put differently, at each selected site, $p\in\partial l\setminus r$, we draw the same random bonds J_{lp} for both populations.

The EA order parameter q_{EA} is then calculated using any of the two replicas a or b, the choice being immaterial as they should produce the same results. For the specific case of $T/J_0=0$, the cavity equations require a modification [158]. In this zero-temperature limit, we have extracted the known values of ρ corresponding to the SG-FSG and FSG-F transition lines from Ref. [159]. These values are $\rho=0.86950(3)$ for the SG-FSG transition and $\rho=0.91665(5)$ for the FSG-F transition.

Finally, the Nishimori line is located at $J_0/T_N = \ln[\rho/(1-\rho)]/2$, independently of the Bethe lattice connectivity, and goes through the multicritical point (ρ_{\star}, T_{SG}) , as expected (see Sec. 3.3). All these results are shown in the phase diagram of Fig. 1.9

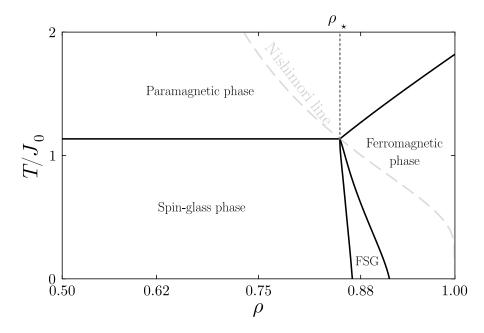


Fig. 1.9: Phase diagram of the RBIM on the Bethe lattice with $\kappa+1=3$ at zero external field (see Sec. 1.5.2). The phase boundaries (solid lines) between the paramagnetic, ferromagnetic, and spin-glass phases, as well as a portion where magnetized spin-glass behavior takes place, denoted here as the ferromagnetic-spin-glass phase (FSG) are obtained as described in the text. For reference, the Nishimori line (gray dashed line) and the multicritical point at $(\rho_\star, T_\star/J_0) = \left(\frac{1+\sqrt{2}}{2\sqrt{2}}, \frac{1}{\operatorname{atanh}(1/\sqrt{2})}\right)$ (gray dotted line) are also included.

1.6 Percolation of the FK-CK clusters

In the following, we will consider the percolation of FK–CK clusters formed by positive (\uparrow) spins—i.e. the equilibrium state is positively magnetized. Before presenting the solution for FK–CK percolation in the frustrated RBIM, let us first examine the case $\rho=1$, which corresponds to the standard ferromagnetic Ising model.

1.6.1 Ising model

Consider the percolation probability $P_o^{(n)}$, defined as the probability that the spin at site o is up, belongs to the percolating cluster, and has n nearest neighbors that are also up, for a generic site o in the lattice. Complementarily, we define the probability $Q_o^{(n)}$, which—unlike in the simple site percolation case—denotes the probability that the spin at site o is up but does not belong to the percolating cluster, while having n neighboring spins that are up. Compactly, we write these conditions as

$$P_o^{(n)} \equiv \Pr \left\{ s_o = \uparrow \land o \in \mathcal{C}_{\infty} \land \sum_{l \in \partial o} \delta_{s_l, \uparrow} = n \right\} ,$$

$$Q_o^{(n)} \equiv \Pr \left\{ s_o = \uparrow \land o \notin \mathcal{C}_{\infty} \land \sum_{l \in \partial o} \delta_{s_l, \uparrow} = n \right\} .$$

$$(1.63)$$

where C_{∞} denotes the set of sites belonging to the infinite percolating cluster. Clearly, n can take values according to $n \in \{0, 1, \dots, c+1\}$. Summing over all possible values of n we obtain the overall probabilities for the central site o to belong and not belong to the

percolating cluster, P_o and Q_o , respectively. These last two probabilities are related to the configuration probability defined in Eq. (1.51) via

$$\eta_o(\uparrow) = P_o + Q_o,\tag{1.64}$$

letting us parametrize all our expressions in terms of the percolation probability, P_o , and the spin up configuration probability $\eta_o(\uparrow)$. For the Ising model $\rho=1$, every site interacts to any other with the constant ferromagnetic coupling J_0 , preserving spatial-translational invariance. Since we are considering a generic site o, we can drop this label and simply write $P^{(n)}$ —the probability that a randomly chosen site belongs to the percolating cluster, with its spin up and the spins of n of its nearest neighbors also up. This leads to the equations defined as

$$P^{(n)} = \frac{1}{\mathcal{Z}_{\text{site}}} e^{\beta h + \beta J_0(2n - \kappa - 1)} \binom{\kappa + 1}{n} (1 - \eta(\uparrow))^{\kappa + 1 - n} \cdots \sum_{l=1}^{n} \binom{n}{l} \pi^l (\eta(\uparrow) - \pi)^{n - l} \sum_{s=1}^{l} \binom{l}{s} p_B^s (1 - p_B)^{l - s} ,$$
(1.65)

where $\mathcal{Z}_{\text{site}}$ is the normalization factor ensuring $P+Q=\eta(\uparrow)$, with

$$P = \sum_{n=0}^{\kappa+1} P^{(n)} , \qquad Q = \sum_{n=0}^{\kappa+1} Q^{(n)} , \qquad (1.66)$$

and π is the *cavity* percolation probability, where the effect of one neighbor (the cavity site) is neglected and thus

$$\pi = \sum_{n=0}^{\kappa} \pi^{(n)} , \qquad (1.67)$$

with the constituent cavity percolation probabilities $\pi^{(n)}$ following the recursive relation

$$\pi^{(n)} = \frac{1}{\mathcal{Z}_{cav}} e^{\beta h + \beta J_0(2n - \kappa)} \binom{\kappa}{n} (1 - \eta(\uparrow))^{\kappa - n} \cdots$$

$$\sum_{l=1}^{n} \binom{n}{l} \pi^l (\eta(\uparrow) - \pi)^{n-l} \sum_{s=1}^{l} \binom{l}{s} p_B^s (1 - p_B)^{l-s} . \tag{1.68}$$

By performing the sum as indicated in Eq. (1.67), we obtain:

$$\pi = \frac{1}{\mathcal{Z}_{cav}} e^{\beta h - \beta J_0 \kappa} \left\{ \left[\left(e^{2\beta J_0} - 1 \right) \eta_{cav}(\uparrow) + 1 \right]^{\kappa} - \left[\left(e^{2\beta J_0} - 1 \right) \eta_{cav}(\uparrow) - p_B \pi e^{2\beta J_0} + 1 \right]^{\kappa} \right\}.$$

$$(1.69)$$

Once again, \mathcal{Z}_{cav} is the normalization factor for the condition $\pi + q = \eta_{\text{cav}}(\uparrow)$ where q and $\eta_{\text{cav}}(\uparrow)$ are the cavity counterparts of the probabilities Q and $\eta(\uparrow)$, respectively.

We can solve Eq. (1.69) iteratively and then compute the quantity P using Eqs. (1.65) and (1.66). However, since we are primarily interested in the critical properties of the clusters—specifically near P=0—and given that $P=0 \Leftrightarrow \pi=0$, the cavity percolation probabilities π serve effectively as order parameters for the percolation transition. Therefore, focusing on π will be sufficient for our analysis. This approach is particularly advantageous in the frustrated case, where the equations for P become significantly more complicated than those for π .

We can analyze the percolation properties of the Ising model as follows: For a given temperature T_K , there exists an associated external magnetic field $h_{\rm ext}^K$ for which the percolation probability starts to deviate from zero, and the FK–CK clusters percolate. Figure 1.10 illustrates the effect for the pure Ising model ($\rho=1$) on a Bethe lattice with $\kappa+1=3$ at $T_K=1.2~T_c$, where T_c the Ising critical temperature given by Eq. 1.57. The onset of percolation then takes place at $h_{\rm ext}^K\simeq 0.1299(1)$. Repeating this procedure for all $T_K\geq T_c$ identifies the points $(h_{\rm ext}^K,T_K)$ that define the Kertész line [160], which identifies the percolation threshold for the FK–CK clusters in presence of an external field.

Alternatively, the external field can be parameterized in terms of the magnetization it induces, and visualize the Kertész line in the (m,T) plane along with the spontaneous magnetization at zero field. The Kertész line separates a non-percolating region, with high temperature and weak spin correlations, from a percolating region at low temperatures, and spins correlated by the external field, resulting in the presence of the spanning cluster. By construction, in the zero external field limit the Kertész line coincides with the Ising critical point, $(h_{\rm ext}^K=0,T_K=T_c)$. Therefore, the thermodynamic and percolating transitions then match.

Note that the Kertész line does not have a direct physical interpretation in terms of clusters, as the CK bonding probability is not strictly valid when $h_{\rm ext} \neq 0$. In fact, there are two symmetric Kertész lines in the phase diagram: one associated with parallel up spins, the other with parallel down spins, corresponding to positive and negative magnetic fields, respectively. These two lines merge at the critical point $(T_c, h_{\rm ext} = 0)$. Formally, it can be shown that in the presence of an external magnetic field, the bonding probability must be modified to account for the field—for instance, by introducing a ghost spin [124, 161]. Nevertheless, tracing the Kertész line is useful for our purposes, as it highlights the fact that it terminates precisely at the critical point. This, in turn, implies that at $h_{\rm ext} = 0$, the clusters defined using the FK–CK bonding probability percolate exactly at criticality.

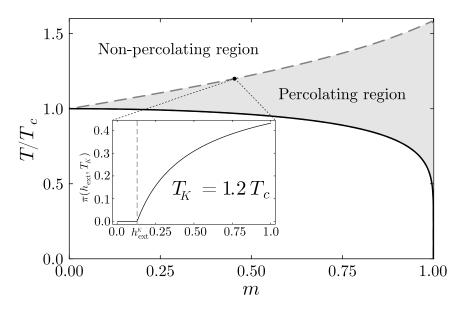


Fig. 1.10: Percolation properties of the pure Ising model ($\rho=1$) on the Bethe lattice with $\kappa+1=3$. Kertész line (dashed line) along with the reversed representation of the magnetization density with (normalized) temperature (solid line). In the Ising model the Kertész line converges to the Ising critical point T_c at $h_{\rm ext}=0$, indicative of the matching between percolation and thermodynamic transitions. The inset shows the cavity percolation probability at $T_K=1.2\,T_c$. The onset $h_{\rm ext}^K=0.1299(1)$, at which the curve detaches from zero, corresponds to a point on the Kertész line.

We shall remember that the percolation probability on the Bethe lattice has critical exponent $\beta_P = 1$ —we have added a subindex P to emphasize that this value is associated to the percolation universality class. In the thermodynamic problem, the order parameter m has two critical exponents near the transition point T_c :

(i) For h = 0 near T_c , the spontaneous magnetization follows the law

$$m \propto |T - T_c|^{\beta},\tag{1.70}$$

with the Ising mean-field critical exponent $\beta = 1/2$.

(ii) At the critical point $T=T_c$ and with a non-zero magnetic field the magnetization density behaves as:

$$m \propto |h_{\text{ext}}|^{1/\delta},$$
 (1.71)

with the Ising mean-field critical exponent $\delta = 3$.

Along the K'ertesz line, the percolation of FK–CK clusters exhibits a crossover. At very high temperatures (corresponding to values close to m=1 in the reversed representation of Fig. 1.10), percolation criticality is governed by the percolation universality class. However, as we approach $T=T_c$, the system crosses over to the Ising universality class.

Since cluster percolation arises due to two control parameters, the temperature T and the external field $h_{\rm ext}$, the percolation exponent β_P associated with the percolation probability P splits into two:

- (i) For $h_{\rm ext}=0$ and temperatures near T_c , we have $P\propto |T-T_c|^{\beta}$, so $\beta_P=\beta$.
- (ii) At $T = T_c$ and for small $h_{\rm ext}$, $P \propto |h_{\rm ext}|^{1/\delta}$, yielding $\beta_P = 1/\delta$.

The Ising universality class is thus recovered at the critical point, as expected. In the following, we will show that this is also the case for the frustrated model, whenever the Ising universality class is present—as in the paramagnetic-to-ferromagnetic portion of the phase diagram in Fig. 1.9.

1.6.2 Frustated RBIM

In this case, the presence of quenched disorder introduces frustration and breaks the spatial-translational invariance of the lattice. As a result, the recursive equations for the percolation probability grow rapidly more complicated with increasing connectivity. Therefore, we here restrict most of our analysis to the case $\kappa+1=3$.

The cavity percolation equations

To construct the percolation probabilities in this case, we build on the above notation. We denote $P_o^{(n)}$ the probability for the site o to belong to the percolating cluster, with n of its nearest-neighbors being spin up, and $Q_o^{(n)}$ the complementary probability that site o does not belong to the percolating cluster, with n of its nearest-neighbors being spin up. As in previous sections, we consider the percolation transition of positive magnetized domains, and hence for the site o to belong to the infinite percolating cluster the spin in o must be up, i.e. $\sigma_i = \uparrow$. Compactly, we write these conditions as

$$P_o^{(n)} \equiv \Pr \left\{ \sigma_o = \uparrow \land o \in \mathcal{C}_\infty \land \sum_{l \in \partial o} \delta_{\sigma_l, \uparrow} = n \right\} ,$$

$$Q_o^{(n)} \equiv \Pr \left\{ \sigma_o = \uparrow \land o \notin \mathcal{C}_\infty \land \sum_{l \in \partial o} \delta_{\sigma_l, \uparrow} = n \right\} .$$

$$(1.72)$$

Clearly, n can take values according to $n \in \{0, 1, ..., c+1\}$. Summing over all possible values of n we obtain the overall probabilities for the site o to belong and not belong to the percolating cluster, P_o and Q_o , respectively. These last two probabilities are related to the configuration probability defined in Eq. (1.51) via

$$\eta_o(\uparrow) = P_o + Q_o,\tag{1.73}$$

letting us parametrize all our expressions in terms of the percolation probability, P_o , and the spin up configuration probability $\eta_o(\uparrow)$. Using the labeling shown in Fig. 1.8, the constituent equations, i.e. those with n=0,1,2,3, are:

$$P_o^{(0)} = 0, (1.74)$$

$$P_o^{(1)} = \frac{e^{\beta h_{\text{ext}}}}{\mathcal{Z}_{\text{site}}} \sum_{(x,y,z) \in C_3} e^{\beta (J_{ix} - J_{iy} - J_{iz})} \pi_{x \to o} p_B^{(ox)} \left[1 - \eta_{y \to o}(\uparrow) \right] \left[1 - \eta_{z \to o}(\uparrow) \right] , \qquad (1.75)$$

$$P_o^{(2)} = \frac{e^{\beta h_{\text{ext}}}}{\mathcal{Z}_{\text{site}}} \sum_{(x,y,z)\in C_3} e^{\beta(J_{ix}+J_{iy}-J_{iz})} [1-\eta_{z\to o}(\uparrow)] \cdots$$

$$\left\{ \pi_{x\to o} p_B^{(ox)} [\eta_{y\to o}(\uparrow)-\pi_{y\to o}] + \pi_{y\to o} p_B^{(oy)} [\eta_{x\to o}(\uparrow)-\pi_{x\to o}] + \pi_{x\to o} \pi_{y\to o} \Phi_{xy}^{(o)} \right\}$$

$$(1.76)$$

$$P_o^{(3)} = \frac{e^{\beta(J_{oi} + J_{oj} + J_{ok} + h_{\text{ext}})}}{\mathcal{Z}_{\text{site}}} \left\{ \sum_{(x,y,z) \in C_3} \pi_{x \to o} \, p_B^{(ox)} [\eta_{y \to o}(\uparrow) - \pi_{y \to o}] [\eta_{z \to o}(\uparrow) - \pi_{z \to o}] \right. \tag{1.77}$$

$$+ \sum_{(x,y,z)\in C_3} \pi_{x\to o} \pi_{y\to o} \left[\eta_{z\to o}(\uparrow) - \pi_{z\to o} \right] \Phi_{xy}^{(o)} + \pi_{i\to o} \pi_{j\to o} \pi_{k\to o} \Psi_{ijk}^{(o)}$$

with

$$\Phi_{xy}^{(o)} = p_B^{(ox)} p_B^{(oy)} + \left(1 - p_B^{(ox)}\right) p_B^{(oy)} + p_B^{(ox)} \left(1 - p_B^{(oy)}\right) , \qquad (1.78)$$

$$\Psi_{ijk}^{(o)} = \sum_{(x,y,z)\in C_3} \left[p_B^{(ox)} \left(1 - p_B^{(oy)} \right) \left(1 - p_B^{(oz)} \right) + p_B^{(oz)} p_B^{(oy)} \left(1 - p_B^{(oz)} \right) \right] + p_B^{(oi)} p_B^{(oj)} p_B^{(ok)} ,$$

$$(1.79)$$

where we used the cyclic permutation group of order 3, $C_3 = \{(i, j, k), (j, k, i), (k, i, j)\}$, to consider compactly all possible combinations of sites having spins up. The overall percolation probability is given by the sum over all constituent contributions

$$P_o = \sum_{n=0}^{c+1} P_o^{(n)}. (1.80)$$

The normalization constant $\mathcal{Z}_{\text{site}}$ enforces the condition $\eta(\uparrow) = P_o + Q_o$. The cavity quantities appearing in Eqs. (1.74)–(1.78), denoted $\pi_{x\to o}$, represent the probability that a site x (= i,j,k) belongs to the percolating cluster in the absence of site o, and its complement, $q_{x\to o}$, corresponds to the probability that x does not belong to the percolating cluster in absence of o. This last quantity does not appear explicitly in the preceding equations because the normalization condition, $\eta_{x\to o}(\uparrow) = \pi_{x\to o} + q_{x\to o}$, has already been imposed.

These equations are constructed considering $\kappa=2$ neighboring spins, reflecting the absence of site o. As with the overall percolation probability, the cavity percolation probability—and its complement—is decomposed into constituent terms corresponding to each possible

number n of neighboring spins in the up state

$$\pi_{x\to o} = \sum_{n=0}^{\kappa} \pi_{x\to o}^{(n)}.$$
 (1.81)

Analogous to the definitions in Eq. (1.72), the constituent cavity percolation probabilities are formally defined as:

$$\pi_{x \to o}^{(n)} \equiv \Pr \left\{ s_x = \uparrow \land x \in \mathcal{C}_{\infty} \land \sum_{l \in \partial x \setminus o} \delta_{\sigma_l, \uparrow} = n \right\},$$

$$q_{x \to o}^{(n)} \equiv \Pr \left\{ s_x = \uparrow \land x \notin \mathcal{C}_{\infty} \land \sum_{l \in \partial x \setminus o} \delta_{\sigma_l, \uparrow} = n \right\},$$

$$(1.82)$$

where in this case n = 0, 1, 2. The explicit form of these equations, following the labeling used in the right-hand side of Fig. 1.8, is given by

$$\pi_{i\to o}^{(0)} = 0 , (1.83)$$

$$\pi_{i\to o}^{(1)} = \frac{e^{\beta h_{\text{ext}}}}{\mathcal{Z}_{\text{cav}}} \left\{ e^{\beta(J_{im} - J_{in})} \, \pi_{m\to i} \, p_B^{(im)} \left[1 - \eta_{n\to i}(\uparrow) \right] + e^{-\beta(J_{im} - J_{in})} \, \pi_{n\to i} \, p_B^{(in)} \left[1 - \eta_{m\to i}(\uparrow) \right] \right\}, \tag{1.84}$$

$$\pi_{i\to o}^{(2)} = \frac{e^{\beta h_{\text{ext}} + \beta(J_{im} + J_{in})}}{\mathcal{Z}_{\text{cav}}} \left\{ \pi_{m\to i} p_B^{(im)} \left[\eta_{n\to i}(\uparrow) - \pi_{n\to i} \right] + \pi_{n\to i} p_B^{(in)} \left[\eta_{m\to i}(\uparrow) - \pi_{m\to i} \right] + \pi_{m\to i} \pi_{n\to i} \left[p_B^{(im)} \left(1 - p_B^{(in)} \right) + p_B^{(in)} \left(1 - p_B^{(im)} \right) + p_B^{(im)} p_B^{(in)} \right] \right\}.$$

$$(1.85)$$

Once where the normalization condition $\pi_{x\to o}+q_{x\to o}=\eta_{x\to o}(\uparrow)$ is imposed by \mathcal{Z}_{cav} . The most explicit form of the cavity percolation probability, using $p_B^{(im)}=1-e^{-2\beta J_{im}}$, is then

$$\pi_{i\to o} = \eta_{i\to o}(\uparrow) \frac{\pi_{n\to i} \,\psi_{in} \left[\eta_{m\to i}(\uparrow) \,\psi_{im} + 1\right] + \pi_{m\to i} \,\psi_{im} \left[\psi_{in} \left(\eta_{n\to i}(\uparrow) - \pi_{n\to i}\right) + 1\right]}{\left[\eta_{m\to i}(\uparrow) \,\psi_{im} + 1\right] \left[\eta_{n\to i}(\uparrow) \,\psi_{in} + 1\right]} . \quad (1.86)$$

with $\psi_{im} = e^{2\beta J_{im}} - 1$. It is easy to see that if $J_{ij} = J_0 > 0 \ \forall (i, j)$ Eq. (1.86) reduces to the Ising case given by Eq. (1.69), where it can be solved by iterations.

However, introducing antiferromagnetic bonds ($J_{im} < 0$) makes the quantity $\pi_{i \to o}$ depend explicitly on the specific realization of J_{im} . Because these bonds are quenched random variables, $\pi_{i \to o}$ becomes a quenched random variable itself. We again solve these cases using the population dynamics algorithm [127], this time the population being the cavity percolation probabilities $\pi_{i \to o}$.

As anticipated above, we omit the computation of the percolation probabilities P_o . The solution of Eq. (1.80) is more complex and is unnecessary to study the critical properties of the system, as $P_0 = 0 \Leftrightarrow \pi_{x \to 0} = 0$. We are interested in the average cavity percolation probability defined as, $\pi \equiv [\pi_{x \to o}]$. We computed it numerically as the average over a population of cavity percolation probabilities, once their distribution reaches a stationary state.

FK-CK clusters in the RBIM

We can now proceed to solve Eq. (1.86) and analyze its behavior across different regimes to identify the onset of percolation. This analysis will focus on the region of the phase diagram where the paramagnetic-to-ferromagnetic phase transition occurs, that is, for $\rho > \rho_{\star}$.

Figure 1.11 shows the Kertész line in the frustrated RBIM for various $\rho \leq 1$. In the (Ising-like) paramagnetic-to-ferromagnetic transition regime, $\rho > \rho_{\star}$, these lines clearly overlap upon approaching the critical temperature, nicely converging towards their respective T_c . The lines, however, shorten as ρ decreases. In addition to their shortening due to the reduced magnetization – following the growing presence of anti-ferromagnetic bonds – we note an unexpected systematic shortening as ρ decreases.

The numerical convergence of the percolation probability computed using the FK–CK cluster definition becomes increasingly unstable as we get closer to the multicritical point $(T_{\star}, \rho_{\star})$ where a different order—the spin-glass phase—emerges. This suggests that the onset of a competing order may hinder the stability of the random-cluster percolation equations.

In the phase diagram at zero external field ($h_{\rm ext}=0$), shown in Fig. 1.9, a magnetized spin-glass (FSG) phase appears for $\rho>\rho_{\star}$. This phase, characterized by a finite magnetization $m\neq 0$ and a nonzero Edwards–Anderson parameter $q_{\rm EA}>0$, persists up to $\rho\simeq 0.916$. The spontaneous magnetization curve at $\rho=0.9$ (Fig. 1.11) clearly display the onset of this FSG phase at very low temperatures. In the presence of an external field, the spin-glass order extends up to the de Almeida–Thouless (AT) line [157, 162, 163], smoothly connecting to the zero-field FSG region that emerges directly along the spontaneous magnetization curve. This spin-glass phase is accompanied by a reentrant behavior in the spontaneous magnetization curve. However, one should keep in mind that our results for m at zero field are obtained within the replica-symmetric (RS) approximation. Since the FSG phase strictly requires the full replica symmetry breaking (RSB) solution, it remains unclear whether this reentrance is a genuine feature of the model or merely an artifact of the RS approximation.

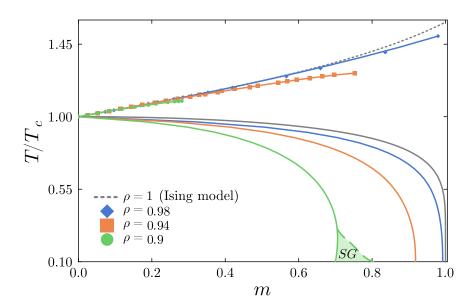


Fig. 1.11: Kertész lines (points; lines are guides for the eye) and spontaneous magnetization (solid lines) for the frustrated RBIM on the Bethe lattice with $\kappa+1=3$ for the temperature in the vertical axis normalized by the critical value at the corresponding ρ , Eq. (1.57). For $\rho=0.9$ and at low temperatures, the spin-glass phase (denoted here with SG) reappears due to the presence of the de Almeida-Thouless (AT) line.

Furthermore, for several values of ρ —down to $\rho=0.86$ lying closely to the multicritical value $\rho_\star\simeq 0.854$ —the FK–CK clusters recover the the mean-field Ising scaling, for the two exponents here considered. This can be seen in Fig. 1.12, where the laws $P\propto |T-T_c|^\beta$ and $P\propto |h_{\rm ext}|^{1/\delta}$ are recovered with $\beta=1/2$ and $\delta=3$. Confirming that the FK–CK clusters belong to the Ising universality class, even in the presence of frustration.

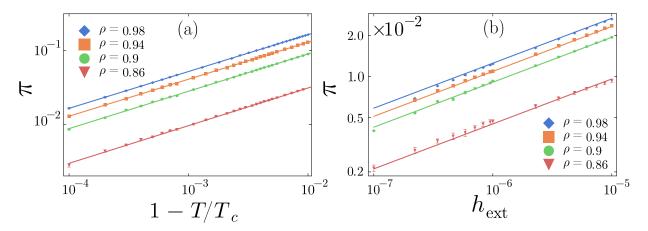


Fig. 1.12: Average cavity percolation probability of FK–CK clusters for the frustrated RBIM as a function of (a) the distance to the critical temperature at zero field and (b) the external field at the critical temperature T_c of each respective ρ . In all cases, upon approaching T_c or $h_{\rm ext}=0$ the critical scaling follows the Ising universality with $\beta=1/2$ and $\delta=3$, respectively.

The α -clusters

Furthermore, as anticipated in Sec. 1.4.2, it is always possible to construct new cluster models parameterized by a quantity α , which can be tuned to define clusters such that the onset of percolation coincides with the critical temperature T_c of the thermodynamic model.

Here, we define these α -clusters through the original proposal of modified FK–CK clusters for frustrated systems [116], already introduced in Sec. 1.4. Specifically, we adjust the bond probabilities on satisfied links—that is, between neighboring sites (ij) where the condition $J_{ij}\sigma_i\sigma_j>0$ holds. For such links, a bond is placed with probability $p_B^{(ij)}=1-e^{-2\beta|J_{ij}|}$. This probability is always strictly positive.

It is well established that with this formulation, the percolation threshold does not coincide with the thermodynamic critical temperature [116]. To address this discrepancy, we introduce a tunable parameter α to enforce this alignment. Consequently, the modified bond probability is defined as

$$p_B^{(ij)}(\alpha) = 1 - e^{-\beta\alpha(J_{ij}\sigma_i\sigma_j + |J_{ij}|)}$$
, (1.87)

with α chosen such that $T_p(\alpha) = T_c$. With this modification Eqs. (1.83)-(1.86) are updated accordingly. We will denote the resulting average cavity percolation probability as $\pi(\alpha)$.

For the dilution $\rho=0.9$, this parameter has been found to be $\alpha\simeq0.7845(7)$. Using this value, we perform the calculation of the critical exponents β and δ , as we did in Fig. 1.12. The result is shown in Fig. 1.13.

Surprisingly, the α -parameter model exhibits the same scaling behavior as the FK–CK clusters. This suggests that, with sufficient numerical precision, it is possible to define positive bond probabilities that can be tuned to yield an effective model with the same critical temperature and critical exponents as the associated thermodynamic transition.

However, the clusters generated by the α -parameter model fail to reproduce the spin–spin correlations required to identify genuine critical clusters. This limitation arises because introducing a non-zero α necessarily modifies the expressions for the correlations $\langle \gamma_{ij} \rangle_W$, introducing an additional term that causes them to deviate from the actual spin–spin correlations $\langle \sigma_i \sigma_j \rangle$. In other words, the Ising critical scaling of the clustering probability is a necessary but insufficient condition for identifying the thermodynamically relevant clusters.

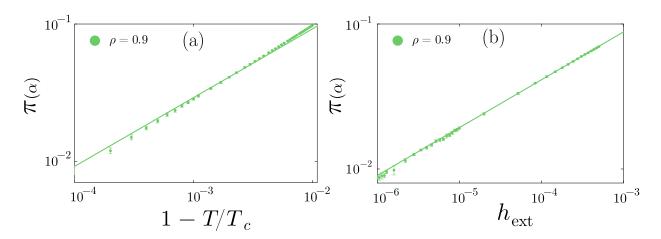


Fig. 1.13: Average cavity percolation probability of α -parameter clusters with $\alpha = 0.7845(7)$ for $\rho = 0.9$ as a function of (a) the distance to the critical temperature at zero field and (b) the external field at the critical temperature T_c . In all cases, upon approaching T_c or $h_{\rm ext} = 0$ the critical scaling follows the Ising universality with $\beta = 1/2$ and $\delta = 3$.

Nevertheless, this may open a new direction to study the spin-glass transition, that has been proven to be even more elusive to be amenable to cluster representations. As the generalized version of the FK–CK clusters are unable to converge passed the multicritical point, we will study a similar α -clusters for the paramagnetic-to-spin-glass transition.

1.7 Spin glass clusters

The instability near the critical dilution ρ_{\star} further confirms the inability of standard FK–CK clusters to properly map the P-SG transition onto a percolating one. This failure is not surprising given that the magnetization is no order parameter for the spin-glass transition and, instead, one needs to turn to overlap variables [35].

Several new definitions of bond probabilities have been proposed [117, 119, 164, 165] in an effort to capture the correlations of the spin-glass order and to provide the basis for cluster-based algorithms aimed at accelerating the dynamics of frustrated systems. However, none of these approaches has achieved universally accepted success [120–122]. Among them, the Houdayer algorithm [118] appears to be the most successful, as it efficiently samples large instances of the two-dimensional Ising spin-glass model on a square lattice. Nevertheless, this method cannot be generalized to higher dimensions.

For d>2, other alternatives yield a percolation transition typically larger than the thermodynamic one, with the closest value reported by the *two-replica FK–CK clusters*. These two-replica FK–CK clusters were initially proposed by Newman and Stein with the aim of providing tools to mathematically show broken spin-flip symmetry in short-range spin glasses at non-zero temperature [166].

1.7.1 Two-replica FK-CK clusters

These clusters follow arise as a simple duplication of the FK–CK construction in two spin configurations with the same realization of the disorder, commonly known as 'replicas', i.e., bonds are occupied independently in the two replicas according to the modified FK–CK probability given by Eq. (1.87). This translates into the bond probability given by

$$p_{ij}(\zeta_{ij}) = \left(1 - e^{-\beta|J_{ij}| - \beta\zeta_{ij}J_{ij}}\right)^2$$
, (1.88)

where

$$\zeta_{ij} = \begin{cases} +1 & \text{for } \sigma_i \sigma_j = 1\\ -1 & \text{for } \sigma_i \sigma_j = -1 \end{cases}$$
 (1.89)

The new percolating equations come with a simplification, as now the sign of the central spin has no implication on its feasibility to belong to the percolating cluster, we can define two general functions for the site o to belong (or not) to \mathcal{C}_{∞} . Using the same labeling as in Fig. 1.8, where the current site is o, and the nearest neighbors are i, j, k (now excluding k as the cavity site), these functions are

$$\mathcal{P}_{\sigma_{i}\sigma_{j}} = \pi_{i\to o} p_{B}^{(oi)}(\zeta_{oi})(1 - \pi_{j\to o}) + \pi_{j\to o} p_{B}^{(oj)}(\zeta_{oj})(1 - \pi_{i\to o})
+ \pi_{i\to o} \pi_{j\to o} \left[p_{B}^{(oi)}(\zeta_{oi})(1 - p_{B}^{(oj)}(\zeta_{oj})) + p_{B}^{(oj)}(\zeta_{oj})(1 - p_{B}^{(oi)}(\zeta_{oi})) + p_{B}^{(oi)}(\zeta_{oi}) p_{B}^{(oj)}(\zeta_{oj}) \right]
\mathcal{Q}_{\sigma_{i}\sigma_{j}} = (1 - \pi_{i\to o})(1 - \pi_{j\to o}) + \pi_{i\to o} \pi_{j\to o} (1 - p_{B}^{(oj)}(\zeta_{oj}))(1 - p_{B}^{(oi)}(\zeta_{oi}))
+ \pi_{i\to o} (1 - p_{B}^{(oi)}(\zeta_{oi}))(1 - \pi_{j\to o}) + \pi_{j\to o} (1 - p_{B}^{(oj)}(\zeta_{oj}))(1 - \pi_{i\to o})$$
(1.91)

These expressions are clearly probabilities, properly normalized i.e. $\mathcal{P}(\sigma_a, \sigma_b) + \mathcal{Q}(\sigma_a, \sigma_b) = 1$. The new cavity probability equations are written in terms of these functions, and are shown below in Eqs. (1.94) for $\kappa = 2$.

below III Eqs. (1.93) - (1.94) for
$$k = 2$$
.

$$\pi_{o \to k}^{(0)} = \frac{(1 - \eta_{i \to o}(\uparrow)) (1 - \eta_{j \to o}(\uparrow))}{\mathcal{Z}_{cav}} \left[\eta_{o \to k}(\uparrow) e^{\beta h_{ext} - \beta(J_{oi} + J_{oj})} \mathcal{P}_{\downarrow \downarrow} + (1 - \eta_{o \to k}(\uparrow)) e^{-\beta h_{ext} + \beta(J_{oi} + J_{oj})} \mathcal{P}_{\uparrow \uparrow} \right], + (1 - \eta_{o \to k}(\uparrow)) e^{\beta (J_{ia} - J_{ib})} \mathcal{P}_{\uparrow \downarrow} + \eta_{j \to o}(\uparrow) (1 - \eta_{j \to o}(\uparrow)) e^{\beta (J_{ia} - J_{ib})} \mathcal{P}_{\downarrow \uparrow} + \eta_{j \to o}(\uparrow) (1 - \eta_{i \to o}(\uparrow)) e^{\beta (J_{ia} - J_{ib})} \mathcal{P}_{\downarrow \uparrow} + \eta_{j \to o}(\uparrow) (1 - \eta_{i \to o}(\uparrow)) e^{\beta (J_{ia} - J_{ib})} \mathcal{P}_{\uparrow \downarrow} \right],$$

$$\pi_{o \to k}^{(2)} = \frac{\eta_{i \to o}(\uparrow) \eta_{j \to o}(\uparrow)}{\mathcal{Z}_{cav}} \left[\eta_{o \to k}(\uparrow) e^{\beta h_{ext} + \beta(J_{ia} + J_{ib})} \mathcal{P}_{\uparrow \uparrow} + (1 - \eta_{o \to k}(\uparrow)) e^{-\beta h_{ext} - \beta(J_{ia} + J_{ib})} \mathcal{P}_{\downarrow} \right].$$

$$(1.92)$$

Their reciprocal counterparts are obtained by replacing $\mathcal{P} \to \mathcal{Q}$ on each of Eqs. (1.93)-(1.94). Once again, the overall cavity probability is given by $\pi_{o \to k} = \sum_{n=0}^{\kappa} \pi_{o \to k}^{(n)}$ and the normalization constant \mathcal{Z}_{cav} ensures $q_{o \to k} + \pi_{o \to k}$. As the up-spin requirement was dropped, the normalization factor \mathcal{Z}_{cav} is now obtained from $\pi_{o \to k} + q_{o \to k} = 1$.

As in previous sections, we calculate $\pi \equiv [\pi_{o \to k}]$ using the population dynamics algorithm for zero external field. Varying the temperature, we find the percolation transition T_p , for a value of $\rho < \rho_\star$ where spin glass order replaces the ferromagnetic one. We obtain a percolation transition of $T_p = 1.3178(5)$ a larger value than $T_{SG} \simeq 1.1346$ as expected.

However, an interesting feature of these clusters within this model is that the this value remains consistent for several values of ρ within the paramagnetic-to-spin-glass transition regime, specifically we measured $\rho=0.5,0.6,0.7,0.8,0.85$.

1.7.2 Multiple-replica FK-CK clusters

Recent works [120, 121] have renewed interest in these clusters by arguing that the bond probability in Eq. (1.88) can be generalized to accommodate an arbitrary number of replicas

$$p_B^{(ij)}(\zeta_{ij}, R) = \left(1 - e^{-\beta|J_{ij}| - \beta\zeta_{ij}J_{ij}}\right)^R , \qquad (1.95)$$

as increasing the number of replicas R might drive the percolation temperature closer to the spin glass transition T_{SG} . This parameter plays a role similar to the α parameter introduced in the previous section. Consequently, one might be tempted to tune R to identify the value that aligns the percolation transition with the thermodynamic transition. We find that $T_p(R) = T_{SG}$ when R = 2.860(8).

We have calculated the cavity percolation probability as a function the normalized temperature T/T_{SG} , for $\kappa=2$ and $\kappa=3$, the results are shown in Fig. 1.14. The universality class is difficult to assess here, as for the spin glass order the critical exponent associated to $q_{EA} \propto |T-T_{SG}|^{\beta}$ in mean-field is $\beta=1$, coinciding with the percolation universality class.

Again, this value appears to be independent of ρ , allowing these 'multiple-replica' clusters to correctly identify the multicritical point at ρ_{\star} . Even when κ is varied, the clusters detect that for $\kappa=3$, the multicritical point given by Eq. (1.59) occurs at $\rho_{\star}\simeq 0.7887$. Consequently, $\rho=0.85$ is no longer a value at which the paramagnetic-to-spin-glass transition is observed, causing the corresponding curve to deviate from the other two at $\rho<\rho_{\star}$. However, for these latter curves $T/T_{SG}\neq 1$, which indicates that the value R=2.860(8) does not correctly tune the critical temperature for this connectivity. This result suggests that the tuned value of R—which enforces the matching between percolation and thermodynamic transitions—depends on the lattice geometry.

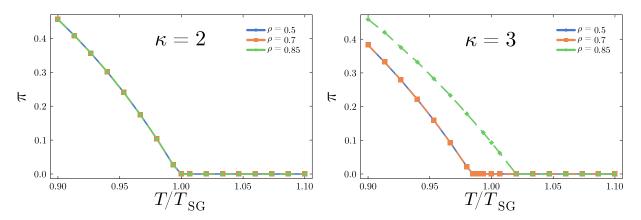


Fig. 1.14: Percolation for the multiple-replica FK-CK with R=2.860(8), for (left) $\kappa=2$ and (right) $\kappa=3$.

Since there is no equivalent proof of FK–CK clusters (or alike) carrying the correct spin glass correlations (either with or without negative probabilities), we cannot determine whether this behavior reflects an underlying physical property of the exponent R that enables the detection of spin glass order, or if it is merely an artificial construction similar to the α -clusters discussed in the previous section. This observation undoubtedly warrants further investigation.

1.8 Conclusions and outlook

In this work, we have generalized the definition of FK–CK random clusters, which capture the physical correlations associated with the ferromagnetic critical point in the simple Ising model, to frustrated models with negative couplings. We first presented a more formal derivation of the cluster construction, extending the proof of Refs. [50, 52] to systems with antiferromagnetic frustration. We then implemented the FK–CK scheme for the frustrated

RBIM, in which the clusters can be constructed by analytically continuing certain bond probabilities to negative values. In particular, we have explicitly verified that, near the line of paramagnetic-ferromagnetic critical points, the critical properties of the clusters fall within the (mean-field) Ising universality class.

This poses a significant obstacle to the idea of constructing an equivalent Swendsen-Wang algorithm for systems with frustration. In such models, the statistical weight of certain cluster configurations can become negative, rendering configurational sampling based on FK–CK clusters unfeasible for accelerating critical dynamics. Although these clusters correctly encode thermodynamic correlations, they cannot be explicitly constructed under frustration.

Our results suggest tha, at least for paramagnetic-to-ferromagnetic transitions, proposing alternative cluster definitions is ultimately futile, since the original FK–CK clusters already capture the correct critical behavior. Furthermore, commonly used indicators of cluster validity, such as the coincidence of thermodynamic and percolation transitions or the appearance of Ising critical exponents, can be misleading. As demonstrated by our construction of the α -clusters, such signatures can be artificially reproduced even when the underlying clusters are not physically meaningful.

This caution also extends to spin glass clusters. Although there is no formal proof that any cluster construction in spin glasses reproduces the relevant thermal correlations, the possibility of artificially defining clusters that appear to coincide with the transition remains. While the detection of the spin glass transition temperature T_{SG} —independently of ρ —by the multiple-replica clusters is indeed intriguing, it should be interpreted with care, as it may represent yet another artificial smoking gun rather than a genuine physical signature.

From a different standpoint, however, this work offers some hope for cluster-based schemes. The fact that the physically relevant clusters cannot be generated through standard constructive schemes does not necessarily mean that they cannot be generated *at all*. Alternative generation methods do remain possible. An appealing prospect entails the use of AI-based generative models or other machine learning-based schemes to *learn rather than construct* clusters. Exactly solvable models on Bethe lattices offer particularly interesting benchmarks for such approaches.

Even in this case, however, there is a—somewhat more subtle—hurdle to consider. Our work demonstrates that critical clusters may only form between parallel spins. However, in systems with antiferromagnetic couplings, configurations with parallel spins connected by such couplings are strongly suppressed by the Boltzmann weight at low temperatures, and hence contribute little to the spin—spin correlation functions. Paradoxically, these same configurations may contribute significantly to the cluster correlation functions. If the bond is present, the cluster weight may be negative; if the bond is absent, the bonding probability may exceed one, thus corresponding to a large weight. These configurations therefore cannot be neglected in the cluster representation. Put differently, correctly generating physical clusters requires sampling rare spin configurations with very low statistical weight.

To illustrate this phenomenon, consider a simple one-dimensional chain of 4 spins with 2 ferromagnetic J>0, and antiferromagnetic, -J, bonds at very low temperature, i.e., near the zero-temperature critical point of the extended chain, with $\beta\to\infty$.

Spin configurations in which the two spins connected by the antiferromagnetic bond are parallel are then strongly suppressed by a factor $e^{-2\beta J}$. Nevertheless, certain cluster configurations associated with these suppressed spin states contribute significantly to the cluster correlation functions, either through a very large (either-negative or positive) statistical weight, as seen for all the configurations with $W_{FK} \to \pm \infty$ as $T \to 0$ in Fig. 1.15.

The relationship between physical and geometrical correlations may therefore be more intricate in the presence of frustration than in clean (unfrustrated) systems. Going back to the model on the Bethe lattice, as the system approaches the spin glass order we observed

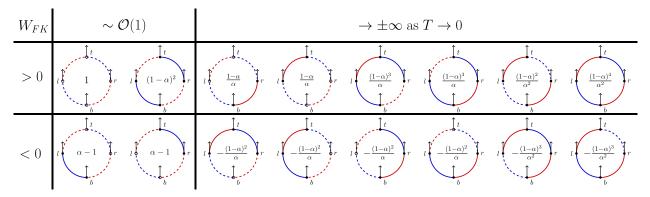


Fig. 1.15: Possible clusters for the spin chain configuration $(\sigma_t, \sigma_r, \sigma_b, \sigma_l) = (\uparrow, \uparrow, \uparrow, \uparrow)$, with periodic boundary conditions, and fixed realization of the random bonds $(J_{tr} = -J, J_{rb} = +J, J_{bl} = -J, J_{lt} = +J)$. The ferromagnetic bonds, J>0, are colored red while the antiferromagnetic bonds, -J, are colored blue. The spins have a filled black circle at their tails if they belong to a cluster, while the circle is left unfilled in the opposite scenario. Despite the associated Boltzmann weight being vanishingly small at low temperatures, the possible clusters reveal non-negligible contributions to the spin-spin correlation functions, coming from the weights of the clusters W_{FK} in the correlated percolation problem. These cluster weights are shown in the center of each respective cluster, where we have set $\alpha=e^{-2\beta J}$. Albeit some contributions are of order one (the four cases shown in the first row), most contributions grow exceedingly large as the temperature decreases (second row). However, there is an equal number of positive and negative contributions, which effectively balance out to yield the physical thermodynamic weight of this fully ordered spin configuration.

that the cavity equations governing percolation probabilities fail to converge. For instance, very close to the multicritical point for the frustrated RBIM, we are unable to find a fixed point of the recursion.

This difficulty also arises far from the critical point. For instance, at higher temperatures and stronger magnetic fields, as evidenced by the shortening of the Kértesz line for values of ρ near the multicritical point ρ_{\star} . In these regions, the system approaches the low-temperature spin-glass phase that emerges along the coexistence line. These observations suggest that the appearance of an additional local minimum in the free-energy landscape – associated with a competing form of order – undermines the stability of the equations governing the percolation of physical clusters. This effect might therefore limit the regime over which relevant clusters can be identified and play a significant physical role.

Chapter 2

SWAP Dynamics for Frustrated Spin Systems

Contents

2.1	Structural glasses and the SWAP method	
	2.1.1	Random First-Order Transition (RFOT) theory 69
	2.1.2	Dynamical facilitation theory
	2.1.3	The SWAP algorithm for structural glasses
2.2		Edwards-Anderson (EA) model
	2.2.1	The model
	2.2.2	Relaxation toward equilibrium
2.3	The Δ	-model for the clean Ising model
2.2 2.3 2.4 2.5 A1		An equivalent Random Bond Ising Model (RBIM)
	2.3.2	Equilibrium properties
	2.3.3	Dynamical properties
2.4	The fr	ustrated Δ-model87
	2.4.1	An equivalent Edwards-Anderson model 89
	2.4.2	The mean-field critical temperature
	2.4.3	Dynamical properties toward the ground state
	2.4.4	Dynamical properties at finite temperature
	2.4.5	Consolidating the key results
2.5	Conclu	usions and outlook
A1		n-and-cut algorithm
A2	More s	snapshots at $L=32$

Introduction

In addition to the well-known problem of critical slowing down near the critical point, glassy systems face another fundamental challenge: their dynamics remain sluggish throughout the entire low-temperature phase. In such cases, novel algorithmic strategies can be developed to enable the efficient simulation of these disordered systems across these regimes.

For spin glasses, methods such as parallel tempering [167, 168], population annealing [169–171], as well as the recently introduced replicated simulated annealing [172, 173], enable the equilibration of larger sample sizes compared to conventional single spin flip Monte Carlo. The special purpose Janus machine has allowed to equilibrate spin glass models of rather large sizes, though only those with discrete variables and couplings [174, 175]. More recently, deep reinforcement learning methods have been explored to find the ground states of finite dimensional spin glasses [176, 177]. Spin glasses are not only interesting physical systems of real magnetic materials, but they also attract theoretical attention because they map to a broad range of hard combinatorial optimization problems.

In the context of structural glasses, Berthier *et al.* [61–63] recently made significant advancements over previous methods [178], achieving accelerated equilibration of several glass-forming liquids. Their breakthrough was first realized in polydisperse mixtures which still exhibit the characteristics of glass formers, and then generalized to many other models [62]. This technique, referred to as SWAP, has proven successful in equilibrating particle systems of unprecedented size, all the way down to the experimental glass transition temperature T_q [179].

Although there is clear empirical evidence for a growing static length scale in glasses, often estimated through point-to-set correlations and related finite-size scaling analyses [180–182], its role in the dynamics remains largely debated [19, 76, 183]. The remarkable acceleration achieved by the SWAP algorithm has been interpreted as direct evidence against static, cooperative explanations of the glass transition, such as those proposed by the random first-order transition (RFOT) theory [76, 183].

The reasoning is that if simply introducing artificial dynamical moves reduces the relaxation time by about 10^{10} , then collective effects—although undoubtedly present—are likely to play only a secondary role in the dramatic slowdown observed as the liquid freezes into a glass state [76]. Instead, this slowdown would be driven primarily by the increasingly constrained dynamics of the system. This interpretation was contested in Ref. [183], where it was argued that the onset of slow relaxation—linked to the emergence of long-lived, large-scale correlated structures arising from the system's collective thermodynamic behavior—can be shifted to lower temperatures by the artificial dynamical moves of SWAP, without altering the thermodynamic properties.

In structural glasses, both static and dynamic effects are present, making it challenging to determine whether the observed acceleration has a thermodynamic origin or is instead driven by purely dynamical constraints [71]. In this context, it would be illuminating to adapt the SWAP algorithm to finite-dimensional spin glasses, where the interplay of frustration and quenched disorder generates correlated structures at low temperatures that are characteristic of spin glass order. Unlike structural glasses, such systems are, by construction, devoid of kinetic constraints [184], providing a cleaner setting in which to probe the source of the algorithm's efficacy.

In this chapter, we adapt the SWAP method for application to finite-dimensional spin glass models. In this reshaping we introduce an auxiliary model, denoted as the Δ -model, in which we assign a length to the spin variables, akin to the role of particle diameters in the original SWAP implementation. The exchange of constituents, in our case, the spins, will effectively mitigate the local energy barriers created by the quenched randomness, allowing

us to explore configurations that would otherwise remain inaccessible.

We have chosen to focus on two-dimensional (2d) problems for which we know the equilibrium phases and ground states. The core of this work is the study of the 2D Edwards-Anderson (EA) model, a random magnet with spin glass properties only at zero temperature, but exceedingly long physical relaxation at low temperatures. This 2D problem is not just a theoretical construct: thin film spin glass materials have regained experimental interest in recent years [185]. Moreover, the ground state configurations can be identified exactly with special algorithms [186, 187]. This gives us a knowledge against which we can confront the performance of our algorithm.

The chapter is organized as follows. In Sec. 2.1, we briefly review the phenomenology of structural glasses, emphasizing the two main competing theories: random first-order transition and dynamical facilitation. In this context, we introduce the SWAP algorithm, and briefly explain the implications of its efficiency in the debate regarding the nature of the glass transition. In Sec. 2.2, we review the known properties of the 2d Edwards-Anderson Ising spin glass. In Sec. 2.3, we study the equilibrium and dynamical properties of the Δ -model with constant ferromagnetic bonds, serving as a testing ground to identify and contrast the effects introduced by spin lengths and non-local exchanges. The core of this work is presented in Sec. 2.4, where we apply the algorithm to the frustrated Δ -model associated with the 2d Edwards-Anderson spin glass, analyzing its dynamics at both zero and finite temperatures, as well as its mean-field behavior. Finally, Sec. 3.8 provides a summary of the main results and outlines potential directions for future research.

2.1 Structural glasses and the SWAP method

Glasses are amorphous solids formed when a liquid is cooled—or compressed—rapidly enough to avoid crystallization. More precisely, a glass is produced by cooling the liquid below its glass transition temperature T_g , fast enough to bypass the first-order transition at the melting temperature $T_m > T_g$ where it turns into a crystal. Importantly, T_g is not a true thermodynamic transition but an empirical quantity: it marks the temperature below which the material becomes so viscous that it effectively stops flowing. In practice, T_g is often taken as the temperature where the shear viscosity reaches 10^{13} Poise, though this operational value of T_g depends on the cooling rate used in the experiment [19].

In 1948, Kauzmann [188] noted that extrapolating the entropy of a supercooled liquid—which decreases as the liquid cools—would, at some low temperature T_K , fall below the entropy of the crystal—an apparent paradox since the crystal is the more ordered phase. This hypothetical Kauzmann temperature T_K (lying below T_g) marks the vanishing of the liquid's configurational entropy $\Sigma(T)$, implying an underlying thermodynamic transition to an 'ideal glass' with a unique ground state. This idea of an emergent amorphous order developing over increasingly large length scales has since been supported by some theoretical [189, 190], experimental [180, 191], and numerical [181, 182, 192–197] studies. However, the existence of this growing amorphous order and the underlying thermodynamic singularity remains under active debate. Specifically, the connection of this thermodynamic picture with the dramatic slowing down of dynamics at low temperatures [76, 183].

Two main theoretical perspectives have emerged. Thermodynamic theories propose that the slowdown arises from an underlying equilibrium phase transition—or at least a rapidly growing static order—while kinetic theories view the glass transition as a purely dynamical phenomenon, with no thermodynamic signature [183]. Several influential frameworks have been developed within these perspectives. Below, we outline two representative approaches: the random first-order transition (RFOT) theory, which attributes glass formation to emergent

static order [198], and dynamical facilitation theory, which explains the transition solely through kinetic effects, without invoking thermodynamic singularities [71, 72].

2.1.1 Random First-Order Transition (RFOT) theory

RFOT theory [67, 198] attempts to explain the phenomenology observed in supercooled liquids through finite-dimensional corrections of mean-field models of spin glasses—such as the random *p*-spin and Potts glass models with infinite-range interactions. In these mean-field models, the sluggish dynamics are directly associated to free-energy barriers that become infinite in the thermodynamic limit, effectively trapping the system in long-lived metastable states [41–43, 199].

Initially formalized by Kirkpatrick, Thirumalai, and Wolynes in the 1980s [64–68], RFOT suggests that barriers between equilibrium states remain finite if the model is considered on a finite-dimensional lattice. Going from one state to another is then possible via the nucleation of a lower free-energy state. The barriers then result from the competition between the free-energy difference and the interfacial free-energy between the various amorphous states. Because the scaling of these two quantities converges in high dimensions, we recover the diverging barrier heights in the infinitely-connected limit, so pure states are then well defined. Fundamentally, it poses that the dramatic slowdown in glass-formers is tied to the existence of a plethora of metastable amorphous states in the free-energy landscape [69].

The theory predicts two distinct regimes. At a first characteristic temperature, denoted T^* , the system begins to develop local rigidity. Below this temperature ($T < T^*$), metastable states emerge that can temporarily 'trap' the system, increasing the relaxation time τ_R . These metastable states are exponentially numerous, yielding a finite configurational entropy $\Sigma(T)$. Because of their large number, the system can still decorrelate over time, behaving as a fluid with a transient rigidity described by nonzero shear modulus. As the temperature is lowered further it reaches the Kauzmann temperature T_K , and the system undergoes a phase transition into an ideal glass phase.

When T is reduced below T^* , the system is described by the 'mosaic picture' [67], which adapts the mean-field notion of metastable states to finite-dimensional systems. In mean-field models, metastable states correspond to distinct minima in the free-energy landscape, and are separated by barriers that scale with system size, making them effectively stable against thermal fluctuations. In contrast, real supercooled liquids have finite-range interactions, so barriers are finite: thermal fluctuations can nucleate rearrangements of finite spatial regions, meaning that metastable states are only well defined locally rather than for the entire system. Different spatial regions of size ξ can therefore be locally close to different metastable minima: within each such region the structure is correlated and resembles one amorphous state, while beyond ξ other regions may correspond to other minima. This results in a 'mosaic' or 'patchwork' structure in real space, where each patch corresponds to a distinct amorphous metastable state.

Interfaces between these patches carry a 'generalized' surface tension Y(T), arising from the mismatch of particle arrangements at their boundaries. The term 'generalized' reflects the fact that the interfacial free-energy cost scales as R^{θ} with $\theta \leq d-1$; when $\theta \neq d-1$, this scaling no longer corresponds to the usual definition of surface tension as energy per unit area. Rearranging a region of size R therefore incurs a free-energy cost proportional to Y(T) R^{θ} , but gains an entropic contribution proportional to -T $\Sigma(T)$ R^{d} due to the large number of available metastable states. Balancing these competing terms yields

the characteristic length scale

$$\xi \sim \left(\frac{Y(T)}{T\Sigma(T)}\right)^{\frac{1}{d-\theta}},$$
 (2.1)

which grows as the configurational entropy decreases, diverging at T_K .

2.1.2 Dynamical facilitation theory

Dynamical facilitation theory proposes that the glass transition is primarily kinetic in origin. In this view, the dramatic slowdown of molecular motion does not arise from an impending thermodynamic singularity; rather, it results from increasingly restrictive rules governing particle dynamics at lower temperatures or higher densities [71, 72]. Within this framework, local regions of enhanced mobility—often called excitations or defects—enable motion in neighboring regions through a process known as facilitation. As a liquid is supercooled, these mobile regions appear less frequently, and structural relaxation occurs only through rare interactions of excitations dynamically facilitating one another. This mechanism naturally leads to dynamical heterogeneity [200–202]. A key feature of supercooled liquids, where one can distinguish highly heterogeneous spatiotemporal patterns of fast and slow spatial regions, in stark contrast to the homogeneous dynamics observed in normal liquids.

To study these ideas in a stripped-down form, kinetically constrained models [73] (KCMs) were introduced. KCMs are minimal lattice models with no interactions—and consequently, trivial thermodynamics—while displaying non-trivial dynamics. The dynamics is introduced by imposing kinetic rules restricting when a region (or site) can relax. For example, in a spin-facilitated model like the Fredrickson-Andersen model [74] or the East model [75], each lattice site has a binary state (e.g. 0 = 'inactive/solid-like' or 1 = 'active/defect'). Spins can flip from 0 to 1 or vice versa only if a certain constraint is satisfied—typically, at least one neighboring site must be active, modeling the idea that mobility can occur only adjacent to other mobile regions. These models are sufficiently tractable for allowing detailed analytical and numerical calculations.

Although KCMs possess no underlying free-energy landscape, they reproduce many aspects of the slow dynamics observed in glass-forming liquids. A superArrhenius temperature dependence of the relaxation time and a consistent picture of the dynamical heterogeneities in a space-time setting are central features that naturally emerge in all these models.

2.1.3 The SWAP algorithm for structural glasses

Recently, Berthier et al. [61–63] demonstrated that introducing non-local Monte Carlo moves, which attempt to exchange the positions of particles with different diameters, can yield an unprecedented acceleration of the relaxation dynamics in supercooled liquids. In particle-based simulations of liquids, the SWAP algorithm introduces an additional fluctuating degree of freedom associated with each particle—its diameter—thereby creating a polydisperse particle mixture, preserving the same properties of glass-forming materials [62, 63, 178, 183].

In the SWAP Monte Carlo algorithm, standard local Metropolis—or molecular dynamics—moves are complemented with these non-local particle exchanges, accepted or rejected in a Monte Carlo step obeying detailed balance. This is pictorially shown in Fig. 2.1, where the two particles of different diameters (green and yellow) are selected and swapped. In systems where the method is efficient, SWAP moves may be rarely accepted; however, when they are, the relaxation dynamics is significantly accelerated. In some cases, the SWAP algorithm can achieve speed-ups of up to 10^{10} compared to traditional methods, enabling

the sampling of equilibrated configurations at very low temperatures that were previously inaccessible.

For poly-disperse particle system

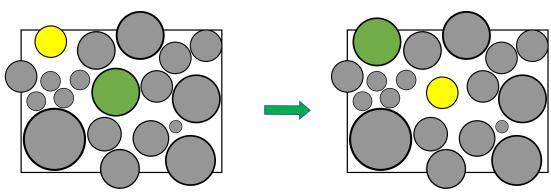


Fig. 2.1: Schematic of the SWAP acceptance move. These moves are alternated with a standard molecular dynamics—or Metropolis—evolution.

This remarkable acceleration was later interpreted as direct evidence against a static, cooperative explanation of the glass transition, such as that proposed by the RFOT theory [76]. The core idea is that switching from standard to SWAP dynamics enhances relaxation only locally, within regions where particle exchanges occur. Consequently, the overall free-energy landscape remains unchanged: as long as both dynamics reach equilibrium, any static (thermodynamic) length scale ξ must be identical for both cases. However, the ratio of equilibration times between SWAP and standard algorithms is enormous—spanning at least four orders of magnitude. This observation suggests that the exponential growth in relaxation times τ_R originates from local barriers that can be dynamically bypassed by SWAP, rather than from a growing static correlation length as proposed in RFOT theory.

A particularly illuminating test of this idea was provided by Gutiérrez, Garrahan, and Jack [77], who implemented SWAP moves within a kinetically constrained model. In their variant of the East model, each lattice site carried not only a binary state (active/inactive) but also an additional local 'softness' variable—analogous to particle diameters in polydisperse liquids. Allowing exchanges of this softness led to a dramatic acceleration of structural relaxation and a marked suppression of dynamical heterogeneity. These results demonstrate that the SWAP method can yield exceptional acceleration even without the presence of a growing static correlation length, consistent with dynamical facilitation theory.

In this work, we present a similar yet complementary study. We adapt the SWAP algorithm to the two-dimensional Edwards–Anderson spin glass model. Unlike kinetically constrained models and glass formers, the Edwards–Anderson system lacks dynamical constraints: its slow dynamics arises from the competition between ferromagnetic and antiferromagnetic interactions introduced by quenched disorder. By applying non-local exchange moves in this setting, we investigate whether the dramatic acceleration observed in other systems persists even when facilitation mechanisms are absent.

2.2 The 2d Edwards-Anderson (EA) model

2.2.1 The model

The Hamiltonian of the model is given by

$$\mathcal{H} = -\sum_{\langle ij \rangle} J_{ij} \sigma_i \sigma_j \;, \quad \sigma_i = \pm 1 \;, \quad i = 1, \dots, N$$
 (2.2)

where $N=L^2$, $\langle ij \rangle$ indicates a sum over nearest neighbors on a two dimensional square lattice with linear size L and periodic boundary conditions. The bond strengths J_{ij} are typically chosen to be discrete, drawn from a bimodal distribution

$$\mathcal{P}(J_{ij}) = \frac{1}{2}\delta(J_{ij} - J) + \frac{1}{2}\delta(J_{ij} + J) , \qquad (2.3)$$

or continuous, drawn from a Gaussian distribution with zero mean, and variance J^2 .

In equilibrium, this 2d EA model has no finite temperature phase transition: it has spin glass ground states and is paramagnetic at non-vanishing temperatures [203–206]. Still, it exhibits a non-trivial slow relaxation toward the equilibrium paramagnetic state at low enough temperatures [60, 186, 187, 192, 207–217].

For Gaussian couplings, the spin glass ground states have a 2-fold degeneracy due to the global spin-reversal symmetry, with a spectrum of excitations growing continuously. For bimodal couplings this degeneracy increases drastically with a gap to the lowest excitations. The Gaussian system undergoes domain growth at T=0 [218]. In the bimodal case, clusters of spins maintain their relative orientation in all ground states and form a backbone on which coarsening takes place while all other spins behave paramagnetically [211]. Recently, the bimodal case was used to study the growth of glassy domains [217], and it was found that the dynamics is controlled by a single timescale that diverges upon approaching T=0 in an Arrhenius-like manner.

2.2.2 Relaxation toward equilibrium

The sluggish dynamics of this system can be assessed from a temperature quench from $T=\infty$ to a temperature $T\gtrsim 0$. As outlined in the introduction, a straightforward way to establish the presence of aging is to calculate the spin auto-correlation function

$$C(t, t_w) = \frac{1}{N} \sum_{i=1}^{N} [\langle \sigma_i(t)\sigma_i(t+t_w)\rangle], \qquad (2.4)$$

where t_w is the waiting time—also referred to as the 'age' of the system—defined as the duration for which the configuration evolves from the initial state before the observation time t begins. The brackets $[\cdots]$ and $\langle \cdots \rangle$ are the standard disorder and thermal averages.

At high temperatures, when t_w exceeds a characteristic time scale τ_R , the auto-correlation function becomes stationary, i.e., $C(t,t_w)=C(|t-t_w|)$. In this regime, all curves with $t_w>\tau_R$ collapse onto a one. This behavior strongly suggests that the system has reached equilibrium: the decay of the correlation function simply reflects the de-correlation between equilibrium configurations. In this situation, one says that aging is interrupted [41–43, 46, 60, 219, 220]. As the temperature decreases, however, the relaxation time τ_R increases significantly, and aging effects reappear, evidenced by the explicit dependence of $C(t,t_w)$ on t_w .

As $T \to 0$ overlap-overlap correlations between spin configurations emerge, and we can study the size of this glassy domains from the 4-point correlation function (or the overlap-overlap correlation function) defined as

$$C_4(r,t) = \frac{1}{N} \sum_{\substack{i,j=1\\ |\vec{r_i} - \vec{r_j}| = r}}^{N} \left[\left\langle \sigma_j^{(1)}(t) \sigma_j^{(2)}(t) \sigma_i^{(1)}(t) \sigma_i^{(2)}(t) \right\rangle \right] , \qquad (2.5)$$

and the characteristic length over which it decays defines the spin glass ordering length—or growing length. It is here estimated from

$$R(T,t) = 2 \int_0^\infty dr \ C_4(r,t) \ ,$$
 (2.6)

as defined in Ref. [60].

The Metropolis dynamics of the 2d EA at $T\simeq 0.5$ yield a super exponential decay, $C_4\sim e^{-(r/R_\sigma(t))^\beta}$, with $\beta>1$. The growing length scales as $R(t)\sim (t/\tau(T))^{1/z}$ with $z\sim 7$ and $\tau(T)$ the Arrhenius characteristic timescale found in Ref. [214]. For low enough temperatures ($T\gtrsim 0$) and bimodal bonds this growing length freezes, saturating at a very short value [221] as seen in Fig. 2.2. The dynamic exponent, plotted in the inset of the same figure, diverges. By increasing the temperature the plateau is surpassed and the evolution persists. The latter growth can be described with the power-law $t^{1/z_{\rm eff}}$, $z_{\rm eff}$ converges to a value close to 8.5 at T larger than 0.3, say.

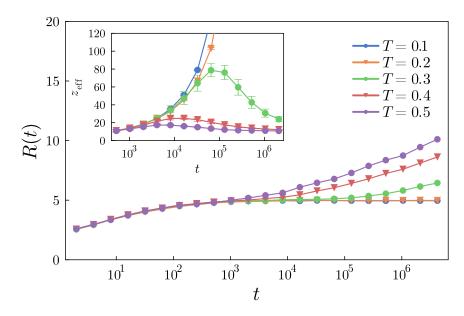


Fig. 2.2: Typical growing length in the $\pm J$ 2d EA model with L=512 evolved with single spin flip dynamics at several temperatures displayed in the key. The inset shows the time dependence of the dynamical exponent $z_{\rm eff}$: it converges to ~ 8.5 after the saturation of the growing length at the value $R_p=4.483\pm0.004$ is superseded at high-enough temperatures (T>0.2).

This saturation is associated the hampering of single spin flips by the presence of frustration, which produces energetic barriers to new favorable configurations—whose effective height increases as the temperature is lowered. We expect that adapting the SWAP method to this problem will accelerate the dynamics by bypassing these local barriers. To this end, we will study a variation of the model in Eq.(2.7) in which the spins have different amplitudes[222]—or 'lengths'—analogous to the softness variable introduced in the kinetically constrained model of Ref. [77]. We will refer to this extension as the Δ -model.

In Fig. 2.3 we provide an schematic example of the expected mechanism whereby the SWAP method accelerates the dynamics. For simplicity, in this sketch we use two spin lengths only, a large one (long arrows and S_{\pm}), and a smaller one (short arrows and s_{\pm}), and bimodal $J_{ij}=\pm J$ interactions. The energy barrier to flip the upper-left highlighted spin is $\Delta E=4J(S^2+Ss)>0$ (as S>0 and s>0); therefore, this spin is blocked at low temperatures. Instead, the energy variation after an exchange of the two highlighted spins is $\Delta E=J(s^2-S^2)<0$ (as s< S). This non-local move will take place since it is energetically favorable, and it may thus help unblocking the upper-left spin and its surroundings. This mechanism will be confirmed by the analysis below.

For a spin lattice system

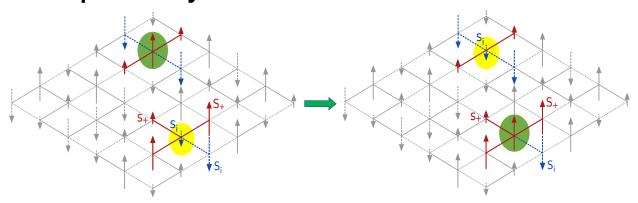


Fig. 2.3: Sketch of a spin configuration of the modified 2d EA model. Two spins are singled out for analysis (surrounded by green bubbles). The neighboring up and down spins are colored red and blue, respectively. The solid (red) and dashed (blue) links represent $J_{ij} > 0$ and $J_{ij} < 0$, respectively. The length of the arrows are proportional to the length of the spins, that is, the local τ_i values. They are here chosen to take only two values, for simplicity.

2.3 The \triangle -model for the clean Ising model

In this section, we introduce the Δ -model for spin systems. We begin by defining the model in the case where the bond strengths are all ferromagnetic, and present the SWAP algorithm adapted to this setting. Concretely, we compare the evolution toward equilibrium of the SWAP algorithm applied to a Δ -model built upon the standard 2d Ising model, to the one of single spin flips. In both cases we work below their finite temperature critical points.

We start from a finite-dimensional 2d Ising Model (IM),

$$\mathcal{H} = -\sum_{\langle ij \rangle} J_{ij} \sigma_i \sigma_j \;, \quad \sigma_i = \pm 1 \;, \quad i = 1, \dots, N$$
 (2.7)

where $N=L^2$, $\langle ij \rangle$ indicates a sum over nearest neighbors on a two dimensional square lattice with linear size L (each pair added once) and periodic boundary conditions. In a clean ferromagnetic model the coupling strengths are all equal $J_{ij}=J>0$.

We will now proceed to 'dress' the Ising spins with a new variable τ_i , creating the Δ -model of interacting soft-spins $s_i = \tau_i \sigma_i$ defined as

$$\mathcal{H} = -J \sum_{\langle ij \rangle} s_i s_j = -J \sum_{\langle ij \rangle} \tau_i \sigma_i \tau_j \sigma_j , \qquad (2.8)$$

with τ_i are independently and initially drawn from a normalized box distribution, $\mathcal{P}_{\tau}(\tau_i)$, i.e.

$$\tau_i \in [1 - \Delta/2, 1 + \Delta/2], \qquad 0 < \Delta < 2.$$
 (2.9)

The average over the τ_i distribution is denoted [...]. The parameter Δ controls the spin length variation: their mean and variance are $[\tau_i] = 1$ and $[\tau_i^2] - [\tau_i]^2 = \Delta^2/12$. The variable Δ has the upper bound $\Delta \leq 2$ to ensure that $\tau_i \geq 0$. The standard Ising model is clearly recovered by setting $\Delta = 0$.

In the rest of this Section we rescale the interactions so as to set J=1. Concretely, the convention is such that the critical temperature of the 2d Ising model is $T_c^{\rm IM}=2.27$.

A ferromagnetic model with spins with variable size has been considered in [222–224]. One of the motivation was to describe 'structurally-disordered magnets' with two (or more) chemically different magnetic components. Specifically, the critical properties of model with variable Ising spin lengths drawn from a bimodal distribution function, was studied in detail in Ref. [224], with special emphasis on its critical properties. We will comment on their findings when discussing the critical properties of our model.

At the initial time of the simulation we need to choose the orientation of the Ising spins σ_i and the lengths τ_i of the spins s_i . For the former we consider two cases, $\sigma_i = \pm 1$ with probability a half, mimicking an infinite temperature initial state, or $\sigma_i = 1$ for all i representing a zero temperature one. For both of them, we draw the τ_i independently from the box distribution in Eq. (2.9).

The numerical simulation is implemented as follows: at each Monte Carlo (MC) sweep, the nature of the microscopic moves is

$$\begin{cases} \text{ with } p_{\text{swap}} & \mapsto N \text{ (non-local) exchange attempts} \\ & (\sigma_i, \tau_i) \leftrightarrow (\sigma_j, \tau_j) \ , \\ \text{with } 1 - p_{\text{swap}} & \mapsto N \text{ single spin flip attempts} \\ & \sigma_i \to -\sigma_i \ . \end{cases}$$

The microscopic moves are accepted with the Metropolis acceptance probability $\mathbb{P}_{\rm acc} = \min(1, e^{-\beta \Delta E})$. The inverse temperature is $\beta = 1/T$ and ΔE is the energy variation due to the *i*-th spin flip or the spin exchange between the *i*-th and *j*-th spins, chosen with the constraint of being more than one lattice spacing apart, making the exchanges strictly non-local. For this model, the energy variation employed in the acceptance probability $p_{\rm acc} = \min(1, e^{-\beta \Delta E})$, is

$$\Delta E = \begin{cases} 2s_i \sum_{j \in \partial i} s_j & \sigma_i \to -\sigma_i \\ (s_i - s_j) \left(\sum_{k \in \partial i} s_k - \sum_{k \in \partial j} s_k \right) & s_i \leftrightarrow s_j \end{cases}$$

The first line corresponds to the flipping of the i-th site and the second one to the exchange between the spins on the i-th and j-th sites, respectively, which are more than one lattice spacing apart.

In later versions of the algorithm, we found that the overall number of accepted moves could be increased by considering pure length exchanges, $\tau_i \leftrightarrow \tau_j$, in the SWAP step. The following results hold for either option, although the latter is recommended.

The symbol ∂i represents the neighbors of the ith spin, i.e. the four nearest neighbors on the square lattice. For $p_{\mathrm{swap}}=0$, all sweeps consist solely of single spin flips, thus the τ_i variables remain quenched throughout the entire MC evolution. Conversely, for $p_{\mathrm{swap}}=1$, all attempted moves are spin-exchanges, allowing the length variables to fluctuate. Any

intermediate value of p_{swap} yields an evolution such that a fraction $1-p_{\text{swap}}$ of the total number of sweeps keeps τ_i unchanged, while for the remaining p_{swap} the τ_i are annealed. We thus refer to the τ_i as partially annealed. Note that the total set of possible lengths $\{\tau_i\}$ remains fixed during a run of the algorithm. In the following, $\langle ... \rangle$ represents an average over thermal MC noise and initial conditions of the $\{\sigma_i\}$ and [...] stands for the average over different realizations of the partially annealed length variables $\{\tau_i\}$.

2.3.1 An equivalent Random Bond Ising Model (RBIM)

After the introduction of $s_i = \tau_i \sigma_i$, random interactions between the Ising spins σ_i emerge. It is straightforward to see that by separating the Ising degrees of freedom and the lengthones, the product of the latter plays the role of positive random bonds in the original Ising Hamiltonian (2.7). An unfrustrated RBIM is then recovered

$$\mathcal{H} = -\sum_{\langle ij \rangle} \mathcal{J}_{ij} \sigma_i \sigma_j \quad \text{with} \quad \mathcal{J}_{ij} = J \tau_i \tau_j ,$$
 (2.10)

that is, a specific and structured distribution of the couplings \mathcal{J}_{ij} induced by the one of the τ_i lengths. This distribution is neither the box nor the bimodal one usually considered in the literature. Being the variables τ_i and τ_j i.i.d. for $i \neq j$, then $\mathcal{P}(\mathcal{J}_{ij}) = \mathcal{P}_{\tau}(\tau_i)\mathcal{P}_{\tau}(\tau_j)$, so that the probability distribution function (pdf) of the new couplings can be found by means of a Mellin transform, yielding

$$\mathcal{P}(\mathcal{J}_{ij}) = \begin{cases} \frac{1}{\Delta^2} \log \left(\frac{\mathcal{J}_{ij}}{u_-^2} \right) & u_-^2 \le \mathcal{J}_{ij} \le u_- u_+ \\ \frac{1}{\Delta^2} \log \left(\frac{u_+^2}{\mathcal{J}_{ij}} \right) & u_- u_+ \le \mathcal{J}_{ij} \le u_+^2 \\ 0 & \text{elsewhere} \end{cases}$$

with $u_{-} = 1 - \Delta/2$ and $u_{+} = 1 + \Delta/2$.

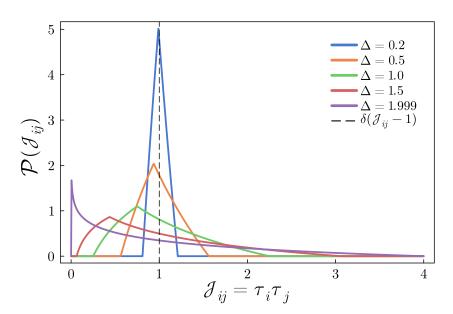


Fig. 2.4: The probability distribution function of the coupling strengths $\mathcal{J}_{ij} = \tau_i \tau_j$ (J = 1) arising from the product of the spin-length variables. Several values of the length-controlling parameter Δ given in the key are considered.

The mean and variance for J=1 are

$$[\mathcal{J}_{ij}] = 1$$
, $[\mathcal{J}_{ij}^2] - [\mathcal{J}_{ij}]^2 = \frac{\Delta^2}{144}(\Delta^2 + 24)$. (2.11)

A plot of $\mathcal{P}(\mathcal{J}_{ij})$ for several values of the length controlling-parameter Δ is shown in Fig. 2.4, where it is clear that as we take $\Delta \to 0$ the pdf tends to a Dirac-delta distribution centered at 1 as it is expected in the Ising spin limit.

These new distributions maintain $\mathcal{J}_{ij} \geq 0$ for all Δ , so frustration is avoided. However, notice that the \mathcal{J}_{ij} are not i.i.d. variables, as for two bonds with a common spin, k, the exchanges are correlated

$$[\mathcal{J}_{ik}\mathcal{J}_{kj}] = [\tau_i \tau_k^2 \tau_j] = [\tau_i][\tau_k^2][\tau_j] = 1 + \frac{\Delta^2}{12}$$

$$\neq [\mathcal{J}_{ik}][\mathcal{J}_{kj}] = [\tau_i][\tau_k]^2[\tau_j] = 1.$$
(2.12)

Therefore, the joint pdf of all couplings \mathcal{J}_{ij} is not just the product of the individual pdfs $\mathcal{P}(\mathcal{J}_{ij})$.

We will characterize the dynamics using both interpretations: firstly, as a ferromagnetic model with soft-spins and secondly, as a RBIM with the above kind of bonds.

2.3.2 Equilibrium properties

In order to appropriately describe the coarsening dynamics, we need to first locate the equilibrium critical temperature. In particular, we have to establish its dependence on the parameter Δ and also characterize the equilibrium properties in the spontaneous symmetry broken phase and close to the critical point. Finally, we have to prove that the equilibrium properties do not depend on the microscopic dynamic rules.

Since the equilibrium configurations below the critical point should be magnetized, for this study we initiate all simulations in σ -ordered configurations, $\sigma_i=1$ for all i. In this way we force a positive magnetization at low temperatures. The length variables τ_i are drawn from the box distribution initially. In this way, we sample different positively magnetized initial configurations of the s_i soft spins.

The averaged magnetizations

Two magnetization densities can be defined,

$$m_s = \frac{1}{N} \sum_{i=1}^{N} \left[\langle s_i \rangle \right], \qquad (2.13)$$

$$m_{\sigma} = \frac{1}{N} \sum_{i=1}^{N} \left[\langle \sigma_i \rangle \right], \qquad (2.14)$$

where the angular brackets denote average over thermal noises and the square brackets average over the distribution of the τ_i s. If one assumes that, in equilibrium, the average $[\langle \tau_i \sigma_i \rangle]$ factorizes as $[\tau_i][\langle \sigma_i \rangle]$, and using the fact that $[\tau_i] = 1$ (for $L \to \infty$), then

$$m_s = \frac{1}{N} \sum_i [\tau_i] [\langle \sigma_i \rangle] = m_\sigma \quad \forall T, \Delta .$$
 (2.15)

This hypothesis is put to the test in Fig. 2.5 where we plot both magnetization densities for an equilibrated L=160 lattice and a not too strong disorder, $\Delta=1$. The data show a very

small systematic deviation, with $m_{\sigma} \leq m_s$, suggesting that there might be a weak correlation between the τ_i and σ_i variables along the evolution, disappearing at high temperatures when the paramagnetic disordered phase prevails with $m_s = m_{\sigma} = 0$.

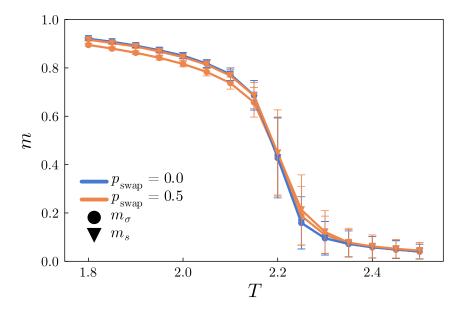


Fig. 2.5: The temperature dependence of both equilibrium magnetization densities (2.13) and (2.14) in the Δ -model with L=160 and $\Delta=1$. A hundred zero temperature, completely ordered, initial conditions with different choices of the τ_i variables were evolved with both single-spin flip dynamics and the SWAP method with $p_{\rm swap}=0.5$ during $t_{\rm max}=2^{16}$ MC-sweeps. After this time, considered to be sufficient for equilibration, we sampled the magnetization densities along 10^4 MC-sweeps. The small difference between the σ and s magnetization densities drops at high temperatures, as expected for the convergence of $m_{\sigma}=m_s$ to zero in the paramagnetic phase. Here and in all other plots the error bars are estimated from the standard deviations.

Indeed, the difference is due to the fact that thermal fluctuations tend to favor the reversal of 'shorter spins', since these flips cost less energy than the ones of longer spins, and hence m_s is slightly higher than m_{σ} in the ordered phase. (The same features are visible within the domains in the out of equilibrium snapshots in Fig. 2.11).

Furthermore, in Fig. 2.5 we also show equilibrium data obtained with the two microscopic dynamics under study, pure single spin-flip and SWAP. The static equilibrium properties are blind to the choices $p_{\rm swap}=0.5$ and $p_{\rm swap}=0$. The equivalence between the datasets built with the two dynamic rules can be verified for other values of $p_{\rm swap}$ in the range $0 \le p_{\rm swap} < 1$ and other values of Δ .

The distribution of the soft-spin s_i values at three times, t=0 and two subsequent ones, for evolutions that led to a positive magnetized domain are shown in Fig. 2.6. The t=0 data are the distribution of the spin-lengths Eq. (2.9), multiplied by ± 1 . The evolution drives the system to positive magnetization and the weight of the pdf progressively moves to the positive support. Note that the two peaks are not symmetric around their midpoints for t>0. Indicating that in the ordered phase, larger spins are more likely to point in the direction of positive magnetization, and the thermal fluctuations are dominated by shorter spins.

The relaxation time depends on the size of the system, temperature and the length-controlling parameter (or disorder-width), $\tau_R(L;T,\Delta)$. Therefore, by increasing L to reduce finite size effects we are in turn increasing the relaxation time, and this makes equilibrium harder to access. In order to test the equilibrium properties at large values of Δ and, in

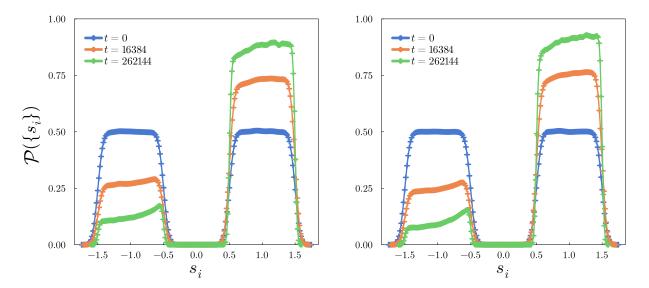


Fig. 2.6: Probability distribution of the spins s_i after a quench to $T=0.77\,T_c$ from an infinite temperature initial condition, with (left) single-spin-flip kinetics and (right) SWAP dynamics ($p_{\text{swap}}=0.5$). Data were sampled using 50 realizations of the τ_i s, with L=128 and $\Delta=1$.

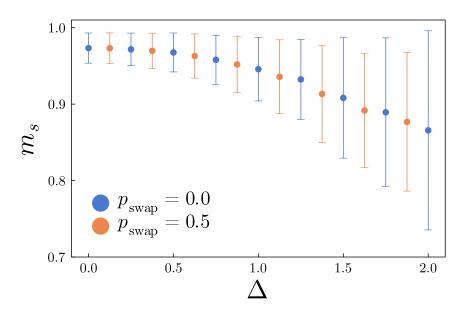


Fig. 2.7: The Δ dependence of the spin s magnetization density for L=16 using both kinds of dynamics. The temperature is set to T=1.67 in all cases which corresponds to $0.74 < T/T_c < 0.84$ depending on the Δ considered. The equilibration time was set to 2^{18} MC-sweeps and data points and error bars were calculated using initial states with $\sigma_i=1$ and 800 choices of the τ_i .

particular, the fact that they do not depend on the microscopic dynamics, it is convenient to use small L. In Fig. 2.7 we compare the Δ dependence of the equilibrium magnetization obtained with single spin flip and SWAP dynamics. Just a few averaged values show a deviation, with the SWAP ones being slightly below the single spin flip ones. This difference is not systematic and in any case very weak so we do not consider it relevant.

The phase transition

A second order ferromagnetic-paramagnetic transition separates magnetized and paramagnetic phases at a Δ dependent T_c . The Binder Cumulant, defined as

$$g_L^s = 1 - \frac{\left[\left\langle \left(\sum_{i=1}^N s_i\right)^4\right\rangle\right]}{3\left[\left\langle \left(\sum_{i=1}^N s_i\right)^2\right\rangle\right]^2},$$
(2.16)

allows one to pin-down the critical point. It locates the critical temperature where the g_L^s data for different L cross, as displayed in Fig. 2.8. We find $T_c(\Delta=0)\sim 2.27$ in good agreement with $T_c^{\rm IM}$. Then, $T_c(\Delta=1)=2.17$, a slightly lower value than $T_c^{\rm IM}$.

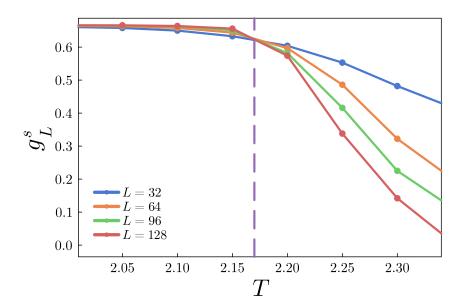


Fig. 2.8: Binder cumulant's temperature dependence for the equilibrium spin s magnetization defined in Eq. (2.16) obtained with single spin flip dynamics and $\Delta=1$. The results are equivalent for equilibrium data obtained with the SWAP evolution with $p_{\rm swap}=0.5$.

In Fig. 2.9 we plot the full Δ dependence of the critical temperatures estimated from the crossing of the Binder parameter. The continuous line is a linear fit which represents the data rather accurately. The range of variation of T_c , 2.05-2.27, with the parameter $\Delta \in [0,2]$ is of the same order as the one found in other works for the conventional RBIM with a box distribution of couplings and $[J_{ij}]=1$ [225]. The analysis of the Binder cumulant of the magnetization m_σ and the two equilibrium magnetizations obtained with the SWAP method yield a $T_c(\Delta)$ which is equivalent to this one within error bars. Interestingly enough, the calculation of the critical temperature in mean-field gives a different behaviour, as $T_c(\Delta)$ grows with Δ in this case. We will show this in latter sections, for the case where disorder is also introduced through the J_{ij} bonds.

With the critical temperature and its dependency on Δ assessed, we proceed to probe the equilibrium behavior around criticality at $T_c(\Delta)$, and compare it to that of the 2d pure ferromagnetic Ising universality class, for which the magnetization and correlation length critical exponents are

$$\beta = 0.125$$
, $\nu = 1$, (2.17)

respectively. In Fig. 2.10 the equilibrium magnetizations m_s of the model with $\Delta=1$ and different system sizes are scaled using the Ising values (2.17). The data fall on a single master curve, confirming that the critical properties of the 2d IM are preserved in the Δ -model.

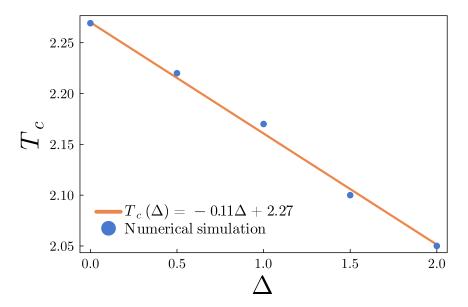


Fig. 2.9: The critical temperature T_c estimated with MC simulation (datapoints) and a linear fit (curve) against Δ . The 2d IM critical value is found at $\Delta=0$. The variation of T_c with Δ is quite weak.

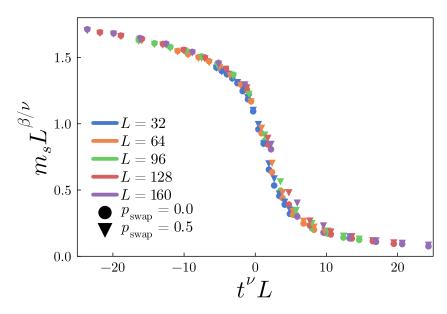


Fig. 2.10: Critical scaling of the equilibrium magnetization density m_s , with $t \equiv (T - T_c)/T_c$ the reduced temperature. $\Delta = 1$. We fixed the critical exponents to the values of the clean ferromagnetic 2d IM, $\beta = 1/8$ and $\nu = 1$, and we used the estimated value of T_c for $\Delta = 1$.

In Ref. [224] the critical properties of this very same model with a bimodal distribution of lengths was studied analytically and numerically. It was shown in this paper that, apart from a special length distribution, in all other cases the criticality is the one of the 2d dilute Ising model. According to the Harris criterium, the 2d dilute Ising model is marginal and the critical exponents (apart from logs) are the same as the ones of the 2d IM.

2.3.3 Dynamical properties

Having checked that the equilibrium properties of the Δ -model do not deviate significantly from the ones of the clean ferromagnetic Ising universality class, we proceed to study the domain growth. We start the dynamics from an infinite temperature initial configuration and we perform an instantaneous sub-critical quench. Studying the coarsening phenomena in this set-up will provide us with a direct survey of the microscopic evolution, hence letting us contrast the efficiency of the SWAP method to the one of the standard single-spin-flip kinetics.

Instantaneous configurations

We first present snapshots of the system domains along the Monte Carlo evolution. As with the magnetization densities, there are two ways to analyze the domains: 1) to consider the s_i spins as a whole, keeping track of both their sign and length, or 2) to isolate the Ising dependence σ_i . Clearly, the former carries more information than the latter.

Snapshots of the s_i spins obtained with single spin flip and $p_{\rm swap}=0.5$ SWAP dynamics, are displayed in Fig. 2.11 top and bottom, respectively. As we lost the binary description of the spin variables, we add a color heat map to distinguish between the different local spin lengths.

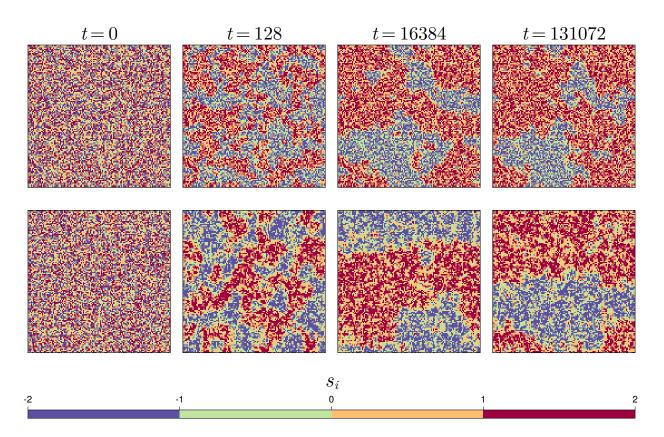


Fig. 2.11: Instantaneous snapshots of the Δ -model with $\Delta=2$ for quench at $T=0.8\,T_c$. First row: single-spin-flip dynamics. Second row: SWAP dynamics. The color scale binned as indicated in the bar shows the lengths of the local spins. Note that the sizes of the domains with the same orientational order look very similar but the darkness within them different. With SWAP longer spins can get together more easily.

As in any coarsening process the size of the domains increases with time. However, here, there are fluctuations within the domains with the same orientation but smaller absolute

value than the rest (green within blue, and yellow within red). The interfaces are typically constituted by a bilayer with short spin length (green and yellow) located in between oriented regions with large spin length (blue and red). The domains built with SWAP (lower row) allow the spins with longer length to get together more easily and hence appear with darker color than the ones grown with single spin flip dynamics.

To quantify the rate at which domain growth occurs we need to measure the time-dependent typical domain size, R(t), along the simulation. There are several numerical ways to get an indirect or direct measurement of this quantity. In the following we extract it from the inverse domain perimeter density and the space-time correlation.

The energy density

The inverse domain perimeter density,

$$R(t) = -\frac{e_{eq}}{e(t) - e_{eq}} = -\frac{e_{eq}}{\delta e(t)}$$
, (2.18)

gives a first estimate of the growing length. The denominator $\delta e(t)$ is the distance between the time-dependent averaged energy density and the equilibrium value, $e_{eq} = \lim_{t \to \infty} e(t)$. For the clean Ising model, $R(t) \sim t^{1/z_d}$ with z_d the dynamical exponent. This in turn fixes a decay exponent for the energy density difference that would go as $\delta e \sim t^{-1/z_d}$. For nonconserved order parameter dynamics $z_d = 2$ while $z_d = 3$ for the locally conserved order parameter ones [226]. Comparing the dynamic exponents should be the simplest way to see whether one algorithm drives the system toward equilibrium faster than the other.

We calculated the equilibrium energy density in the long-time limit ($t>2^{18}$ MC-sweeps) of runs initiated in σ -ordered initial conditions and averaged over 100 realizations of the τ_i s. The time-dependent $e(t) = [\langle \mathcal{H}(t) \rangle]/N$ was computed, instead, after quenches from completely disordered initial conditions, with parameters such that the equilibrium state is ordered. The decay of the excess energy curves is shown in Fig. 2.12 in double logarithmic scale. A power law fit of the energy decays obtained with both methods leads to a time-dependent effective exponent $z_{\rm eff}$ reported in the caption.

After a short transient in which the effective exponent of the single spin flip dynamics is close to 2, the decay of the excess energy slows down and the effective exponent increases in time reaching a value close to 3 in the considered time window (blue data points in the main panel and the inset). The reason for this is that the variable length model has, from the point of view of the Ising spin variables σ_i , quenched random bonds. Therefore, the single spin flip dynamics feels these randomness and its evolution is slowed down, with the effective exponent developing a dependence on the disordered strength, as pointed out in the literature [225, 227–233].

The SWAP method circumvents the slowing down introduced by the variable length of the spins. The dynamical exponent of single-spin-flip kinetics of the 2d IM [226], $z_d = 2$, is recovered when one adds spin exchanges (orange data points in the main panel and the inset), effectively annealing the disorder.

The space-time correlation

We now focus on the one-time spatial correlation function which also carries information on the typical domain size along the evolution. As with the magnetization densities, we either

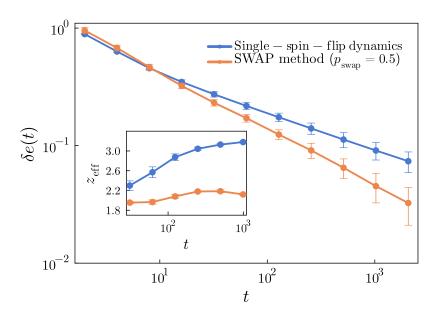


Fig. 2.12: Averaged energy density relaxation of the ferromagnetic Δ -model in double logarithmic scale. Disordered initial configurations prepared at $T_0=\infty$ were evolved at $0.55\,T_c$ using single-spin flip and SWAP kinetics with $p_{\rm swap}=0.5$. Data correspond to $\Delta=1$ and L=128, and were averaged over 100 different initial conditions for both the σ_i s and the τ_i s. A power law fit over moving time windows including three data points yields the effective exponent $z_{\rm eff}$ reported in the inset.

measure the space-time correlation of the complete spin-variables,

$$C_{s}(r,t) = \frac{1}{N} \sum_{j=1}^{N} \sum_{k=1}^{N} \left[\langle s_{j}(t)s_{k}(t) \rangle \right]_{|\vec{r}_{j} - \vec{r}_{k}| = r}$$

$$= \frac{1}{N} \sum_{j=1}^{N} \sum_{k=1}^{N} \left[\langle \tau_{j}\sigma_{j}(t)\tau_{k}\sigma_{k}(t) \rangle \right]_{|\vec{r}_{j} - \vec{r}_{k}| = r}$$

$$(2.19)$$

where the double sum over indices j and k is restricted by the condition on the distance between the spins. We normalize by the number of terms considered. Or else, we measure the space-time correlations of the Ising variables, in direct correspondence with the RBIM interpretation

$$C_{\sigma}(r,t) = \frac{1}{N} \sum_{j=1}^{N} \sum_{k=1}^{N} \left[\langle \sigma_j(t) \sigma_k(t) \rangle \right] \Big|_{|\vec{r}_j - \vec{r}_k| = r} . \tag{2.20}$$

In both cases, $\langle \cdots \rangle$ is the average over σ_i initial conditions and noise realizations of the dynamics on the one hand, and over the 2D equidistant lattice sites on the other. As stated in prior Sections, $[\cdots]$ represents a *disorder average* over several realizations of the lengths $\{\tau_i\}$. We note that

$$C_s(r=0,t) = \frac{1}{N} \sum_i \left[\tau_j^2\right] = 1 + \frac{\Delta^2}{12} ,$$
 (2.21)

$$C_{\sigma}(r=0,t) = 1,$$
 (2.22)

at all times t.

Dynamic scaling states that a *single* domain length R(t) should scale all correlation functions [199]. We now study separately the correlations of the s_i and σ_i spins to confirm that this is indeed the case in this problem.

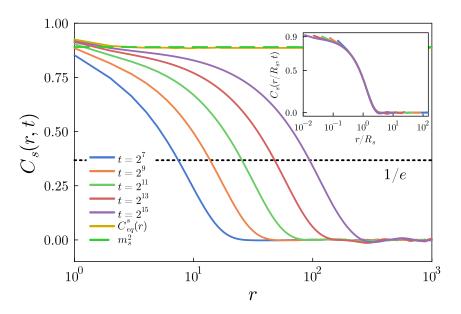


Fig. 2.13: SWAP dynamics with $p_{\rm swap}=0.5$ of an L=2048 ferromagnetic system with $\Delta=1$, quenched from $T_0 = \infty$ to $T \approx 0.77 T_c$. Main panel: space-time correlation of the s spins at several times given in the key. The solid (gold) line is the equilibrium correlation while the dashed green line is m_s^2 in equilibrium. Inset: test of dynamic scaling with the typical domain size $R_s(t)$ estimated from $C_s(R_s(t),t) = C_s(r=0,t)/e$, with $C_s(r=0,t) = 1 + \Delta^2/12$. Averages are performed over 10 runs.

Not too close to the critical temperature, where the equilibrium magnetizations are not too low, one can estimate the typical domain size $R_{s,\sigma}(t)$ from the r such that the correlations $C_{s,\sigma}(r,t)$ decay to, say, 1/e of their zero distance values. A dynamic scaling regime, in which

$$C_{s,\sigma}(r,t) \sim f\left(\frac{r}{R_{s,\sigma}(t)}\right)$$
 (2.23)

is expected for $\xi_{\rm eq} \ll r \ll L$ with $\xi_{\rm eq}$ the equilibrium correlation length and L the linear system size. Numerically, it is convenient to measure r along any of the two axes of the 2D lattice with PBCs.

The decay and scaling of the space-time correlation $C_s(r,t)$ in a model with $\Delta=1$ are studied in Fig. 2.13, using the SWAP method with $p_{\text{swap}} = 0.5$ and single spin-flips. At different times, the curves differ, demonstrating, once again, the out of equilibrium character of the dynamics. The golden curve, which approaches at far distances a finite constant, corresponds to the correlation decay of an ordered initial configuration that we heated up to the desired temperature $T < T_c$ and evolves in equilibrium. At very large r, the two spins in Eqs. (2.19) and (2.20) are expected to become independent and the average factorize, $\langle s_j s_k \rangle \sim \langle s_j \rangle \langle s_k \rangle$, leading to

$$\lim_{r \to \infty} C_{\text{eq}}^s(r) = [\langle s_i \rangle^2] = m_s^2 \,, \tag{2.24}$$

$$\lim_{r \to \infty} C_{\text{eq}}^{s}(r) = [\langle s_i \rangle^2] = m_s^2 ,$$

$$\lim_{r \to \infty} C_{\text{eq}}^{\sigma}(r) = [\langle \sigma_i \rangle^2] = m_\sigma^2 .$$
(2.24)

Moreover, in equilibrium the two magnetization densities are almost identical. The first of these limits is verified numerically in Fig. 2.13 and the other one as well. Dynamic scaling is checked in the inset of Fig. 2.13 and it works equally fine for the σ correlations, see Fig. 2.14. The dynamical scaling master curves, f(x), for C not too close to zero coincide, as can be seen in Fig. 2.14. There are differences when the Cs get close to zero, with oscillations for SWAP which are absent for single spin flip (somehow reminiscent of the oscillations also present when working with local-spin-exhanges as in the case of Kawasaki dynamics [48]).

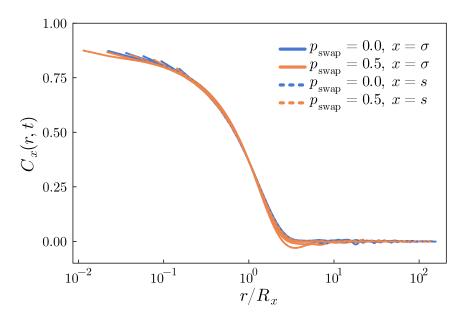


Fig. 2.14: Scaling of the space-time correlations C_s and C_σ for the single-spin flip ($p_{\text{swap}} = 0$) and SWAP dynamics with $p_{\text{swap}} = 0.5$. The same master curve describes both sets of data. Same parameters as in Fig. 2.13.

The growing length for the soft spins

The growing length R_s —measured from the soft-spins space-time correlations—generated with the two updating rules are studied in Fig. 2.15 and Fig. 2.16. In the former, the clean ferromagnetic Ising model's $R_s = R$ is analyzed as a benchmark. In the latter, the R_s of the soft spin model is studied. As stated before, the typical domain sizes measured are fitted via a power law, with an effective exponent which depends on the width Δ ,

$$R_s(t) = \lambda t^{z_{\text{eff}}^{-1}(\Delta)} . \tag{2.26}$$

 λ is a non-universal parameter. The exponent $z_{\rm eff}$ is measured by taking averages over successive time-windows along the domain growth, as we did in the previous Section when the inverse perimeter density was calculated. The convergence of $z_{\rm eff}$ toward $z_d=2$ in the clean case is verified in Fig. 2.15 for single-spin-flip kinetics, while the SWAP method is unable to accelerate the dynamics, and produces a slightly slower convergence.

In Fig. 2.16 we study $z_{\rm eff}$ for the soft-spin model with $\Delta=1$. When the system is evolved with single-spin-flip kinetics, we get $z_{\rm eff}\sim3$, and these update rules are not convenient. However, when we implement the SWAP method, the effective exponent decreases to 2.125, a value that is very close to the theoretical $z_d=2$ of the regular 2d IM with non-conserved order parameter dynamics.

The growing length for the Ising spins

Now, we repeat the analysis above but this time measuring $R_{\sigma}(t)$. Tracking the domain growth of the Ising variables in Fig. 2.17, similar conclusions are reached. The dynamic exponent remains close to 3 for the single-spin-flip updates, while $z_{\rm eff}$ decreases when SWAP is performed, approaching a value close to 2.25 in the numerical interval explored, which is slightly larger than $z_d=2$ for the 2d IM with non-conserved order parameter [226].

Finally, we study the disorder dependence, Δ , of the asymptotic value of $z_{\rm eff}$, which we call $z_{\rm eff}^{\infty}$. We measure it in the last available time-interval. The results are plotted in Fig. 2.18. There is a large increase of $z_{\rm eff}^{\infty}$ with the width of the spin length distribution for the single spin flip dynamics, while there is none, apart from noise, in the SWAP simulations.

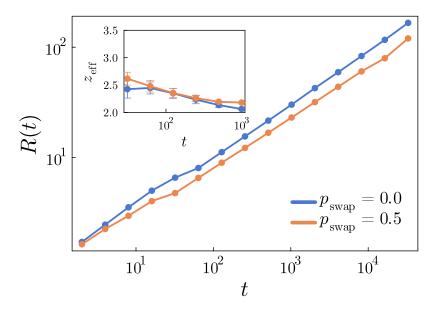


Fig. 2.15: The growing length of the clean Ising model, $\Delta = 0$, with single spin flip and SWAP dynamics. The inset shows the time dependence of the effective exponent z_{eff} , measured by performing a fit of the data at (moving) six consecutive times.

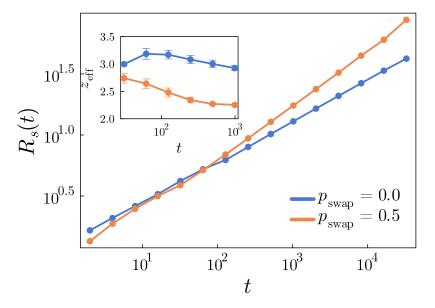


Fig. 2.16: The growing length R_s (estimated from the analysis of the s spins space-time correlation) of the $\Delta=1$ -model, with single spin flip and SWAP dynamics. Sub-critical quench to $T\approx 0.74~T_c$ and L=2048. In the inset, the time dependence of the effective exponent $z_{\rm eff}$, measured by performing a fit of the data at (moving) six consecutive times.

2.4 The frustrated \triangle -model

In this Section, we will restore the site dependence of the original quenched bonds J_{ij} and construct the Δ -model over this case. Specifically, we will consider the 2d Edwards-Anderson (2d EA) model, as introduced in Eq. (2.2) choosing the bond strengths to be drawn from a bimodal symmetric distribution

$$\mathcal{P}(J_{ij}) = \frac{1}{2}\delta(J_{ij} - J) + \frac{1}{2}\delta(J_{ij} + J) . \tag{2.27}$$

Once again, we will set J = 1 in the numerical applications.

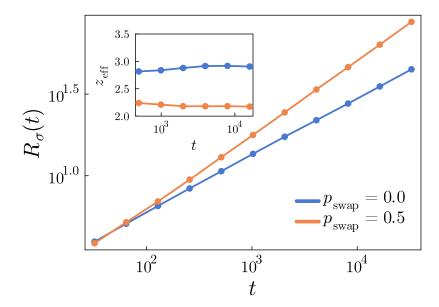


Fig. 2.17: The growing length R_{σ} (estimated from the analysis of the Ising σ spins space-time correlation). Sub-critical quench to $T\approx 0.74~T_c$ and L=2048. The inset shows the effective exponent $z_{\rm eff}$ variation in time, measured by performing a fit of the data at (moving) six consecutive times.

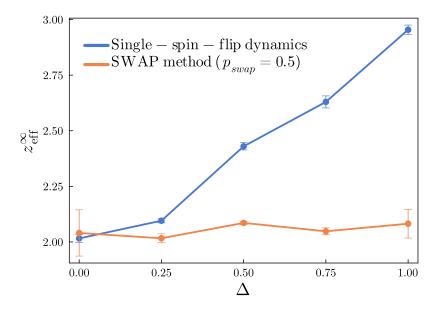


Fig. 2.18: The asymptotic dynamical exponent, $z_{\rm eff}^{\infty}$, estimated from the effective exponent $z_{\rm eff}$ of the spins s growing length, in the latest time interval accessed by the simulation, against the disorder width Δ . Sub-critical quench to $T\approx 0.74\,T_c$, and L=2048.

The Hamiltonian of the frustrated Δ -model takes the form

$$\mathcal{H} = -\sum_{\langle ij\rangle} J_{ij} s_i s_j = -\sum_{\langle ij\rangle} J_{ij} \tau_i \tau_j \sigma_i \sigma_j . \qquad (2.28)$$

Again for $\Delta = 0$, $\tau_i = 1 \ \forall i$, it boils down to the 2d EA model with the bimodal couplings in Eq. (2.2).

2.4.1 An equivalent Edwards-Anderson model

By following the same recipe used to analyze the ferromagnetic model, one can re-express the Hamiltonian with randomly chosen lengths τ_i as an EA Ising spin glass

$$\mathcal{H} = -\sum_{\langle ij\rangle} \mathcal{J}_{ij} \sigma_i \sigma_j , \qquad (2.29)$$

with an uncommon kind of couplings, defined as $\mathcal{J}_{ij} = J_{ij}\tau_i\tau_j$. The \mathcal{J}_{ij} are correlated through the site dependence of the $\{\tau_i\}$, similarly to what happened with the ones of the ferromagnetic model, see Eq. (2.12).

The mean and variance of the new effective couplings \mathcal{J}_{ij} are

$$[\mathcal{J}_{ij}] = 0$$
, $[\mathcal{J}_{ij}^2] - [\mathcal{J}_{ij}]^2 = J^2 \left(1 + \frac{\Delta^2}{12}\right)^2$. (2.30)

Moreover, there is a persistent quenched randomness, the J_{ij} , that is unaffected by the choice of the dynamics, unlike the $\{\tau_i\}$ that remain locally unmodified only when the dynamics do not involve spin exchanges. This randomness is the only source of frustration in the model, as the spin lengths inducing the new effective bonds \mathcal{J}_{ij} are all positive, and do not remove it, independent of the dynamics.

In Fig. 2.19 we show the distribution of the couplings \mathcal{J}_{ij} for three values of Δ . The usual Gaussian distribution with zero mean and unit variance is also shown for reference. The distribution of the couplings \mathcal{J}_{ij} is just two delta peaks at $\pm J$ for $\Delta \to 0$ and it progressively shrinks the gap for increasing Δ until closing it completely when $\Delta = 2$.

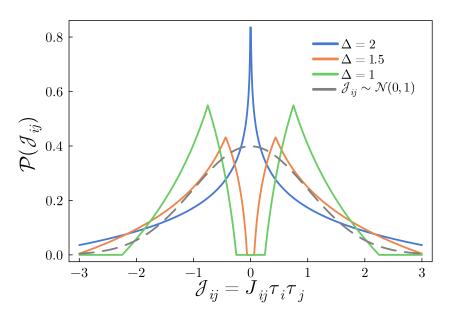


Fig. 2.19: The probability distribution function of the coupling strengths $\mathcal{J}_{ij} = J_{ij}\tau_i\tau_j$, arising from the product of the spin-lengths $\tau_i\tau_j$ and the quenched couplings J_{ij} taking ± 1 values with probability a half. The dashed line represents a Gaussian distribution in normal form for comparison.

2.4.2 The mean-field critical temperature

As anticipated in the previous section, we will calculate the critical temperature of the frustrated RBIM with *fixed* couplings \mathcal{J}_{ij} (i.e. when both the J_{ij} 's and the τ_i 's are quenched) using

two different mean-field methods: (i) the Thouless-Anderson-Palmer (TAP) approach [35, 40] and (ii) the model on the Bethe lattice [126, 127]. The former will yield an equivalent result to the replica-symmetric calculation of $T_{\rm SG}$ in the fully-connected case—i.e. the Sherrington-Kirkpatrick model with \mathcal{J}_{ij} bonds.

(i) TAP approach: After having conveniently scaled the \mathcal{J}_{ij} with N to ensure a the extensive property of observables in the thermodynamic limit, one derives the following equations for the local magnetizations:

$$m_i = \tanh \left[\beta \sum_{i \in \partial i} \mathcal{J}_{ij} m_j + \beta h_i^{\text{ext}} - \beta^2 m_i \sum_{i \in \partial i} \mathcal{J}_{ij}^2 (1 - m_j^2) \right] . \tag{2.31}$$

Assuming a continuous phase transition and taking $h_i^{\text{ext}} \sim 0$ as well, the local magnetization should be $m_i \sim 0$. If, moreover, one replaces \mathcal{J}_{ij}^2 by $[\mathcal{J}_{ij}^2]$ in the Onsager reaction term,

$$m_{i} \sim \beta \sum_{j \in \partial i} \mathcal{J}_{ij} m_{j} + \beta h_{i}^{\text{ext}} - \beta^{2} m_{i} \sum_{j \in \partial i} [\mathcal{J}_{ij}^{2}]$$

$$\sim \beta \sum_{j \in \partial i} \mathcal{J}_{ij} m_{j} + \beta h_{i}^{\text{ext}} - \beta^{2} J^{2} \left(1 + \frac{\Delta^{2}}{12}\right)^{2} m_{i} .$$

$$(2.32)$$

This equation can now be taken to the basis of eigenvectors of the matrix with elements \mathcal{J}_{ij} . Calling \boldsymbol{v}_{μ} the eigenvector associated to the eigenvalue λ_{μ} , and $m_{\mu} = \boldsymbol{m} \cdot \boldsymbol{v}_{\mu}$,

$$m_{\mu} \sim \beta \lambda_{\mu} m_{\mu} + \beta h_{\mu}^{\text{ext}} - \beta^2 J^2 \left(1 + \frac{\Delta^2}{12} \right)^2 m_{\mu} .$$
 (2.33)

The linear susceptibilities are

$$\chi_{\mu} = \frac{\partial m_{\mu}}{\partial h_{\mu}^{\text{ext}}} \bigg|_{\boldsymbol{h}^{\text{ext}} = \boldsymbol{0}} \sim \frac{\beta}{1 - \beta \lambda_{\mu} + \beta^{2} J^{2} \left(1 + \frac{\Delta^{2}}{12} \right)^{2}}.$$
 (2.34)

The first susceptibility to diverge is the one associated to the largest eigenvalue λ_{\max} and this arises at

$$\beta_{\text{SG}} = \frac{\lambda_{\text{max}} \pm \left[\lambda_{\text{max}}^2 - 4J^2 \left(1 + \frac{\Delta^2}{12}\right)^2\right]^{1/2}}{2J^2 \left(1 + \frac{\Delta^2}{12}\right)^2} \,. \tag{2.35}$$

In the Sherrington-Kirkpatrick model, the mean-field limit of the EA model, $\Delta \to 0$ and $\lambda_{\max} = 2J$. Then, $\beta_{\rm SG} = J^{-1}$. If, in the Δ -model, $\lambda_{\max} = 2[\mathcal{J}_{ij}^2]^{1/2}$, which seems reasonable, then

$$\beta_{\rm SG} \propto [\mathcal{J}_{ij}^2]^{-1/2} \quad \Rightarrow \quad T_{\rm SG} = J\left(1 + \frac{\Delta^2}{12}\right) \ .$$
 (2.36)

In Fig. 2.20 we plot, with orange data points, the critical temperature $T_c(\Delta)$ obtained from diagonalizing symmetric matrices with such elements and linear size L=32 (J=1). The data points are consistent with the quadratic dependence on Δ linking $T_{\rm SG}=J$ at $\Delta=0$ and $T_{\rm SG}=4/3\,J$ at $\Delta=2$.

(ii) Bethe lattice: An alternative way to estimate the Δ dependence of the critical temperature with a mean-field approach is to place the model on a Bethe lattice with

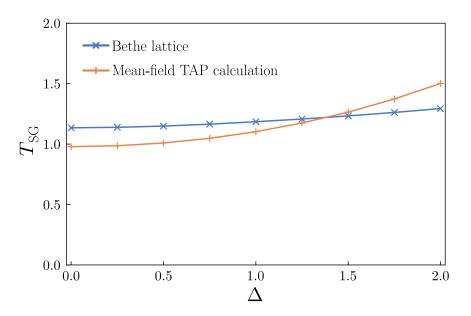


Fig. 2.20: The Δ dependence of the critical temperature in mean-field. Orange data: results for the fully connected model with the TAP approach. Blue data: results for the Bethe lattice model with connectivity $\kappa+1=3$. The trend is the same and the range of variation of the critical temperature with Δ is very weak in both cases.

connectivity $\kappa + 1$. By defining the cavity field acting on site i produced by the effect of $\kappa + 1$ neighboring spins (in the absence of its j-th neighbor), we obtain the recursive equations

$$h_{i\to j}(\tau_i) = \sum_{m\in\partial i/j} \frac{1}{\beta} \operatorname{atanh} \left[\tanh \left(\beta J_{im} \tau_i \tau_m\right) \tanh \left(\beta h_{m\to i}(\tau_m)\right) \right]. \tag{2.37}$$

These equations admit $h_{i\to j}=0$ as a solution on all sites, corresponding to the paramagnetic phase. Since we will be interested in the temperature regime in the vicinity of the critical point, $T \lesssim T_{\rm SG}$, in which the cavity fields are small, we can expand the right-hand-side of Eq. (2.37):

$$h_{i\to j}(\tau_i) \simeq \sum_{m\in\partial i/j} \tanh\left(\beta J_{im}\tau_i\tau_m\right) h_{m\to i}(\tau_m) \ .$$
 (2.38)

Due to the randomness of the spin-amplitudes, these equations must be interpreted as a self-consistent integral equation for the probability distribution of the local cavity fields

$$\mathcal{P}(h|\tau) = \int \prod_{m=1}^{\kappa} \left[\sum_{J_m}^* d\tau_m dh_m \, \mathcal{P}_{\tau}(\tau_m) \, \mathcal{P}(h_m|\tau_m) \right] \delta(h - \tilde{h}), \tag{2.39}$$

with $\tilde{h} = \sum_m \tanh \left(\beta J_m \tau \tau_m\right) h_m$ and the sum \sum^* indicates that there is an implicit 1/2 factor. Using the integral representation of the δ -function, one obtains the following equation for the Fourier transform of the probability distribution

$$\hat{\mathcal{P}}(q|\tau) = \left[\int \sum_{J}^{*} d\tau' p_{\tau}(\tau') \,\hat{\mathcal{P}} \left(q \tanh \left(\beta J \tau \tau' \right) | \tau' \right) \right]^{\kappa} , \qquad (2.40)$$

Assuming that the fields follow a Gaussian distribution with variance, $\sigma_h^2 = \overline{h(\tau)^2} - \overline{h(\tau)}^2$, we have

$$\hat{P}(q|\tau) = 1 - iq\overline{h(\tau)} - \frac{q^2}{2}\overline{h^2(\tau)} + \dots$$

plugging this expression into the equation above, using the fact that $\overline{h(\tau)^n} \ll 1$ for $T \lesssim T_{\rm SG}$, and expanding up to second order one finds

$$\overline{h(\tau)} = \kappa \int \sum_{J}^{*} d\tau' \ p_{\tau}(\tau') \ \tanh(\beta J \tau \tau') \ \overline{h(\tau')}, \tag{2.41}$$

$$\overline{h^2(\tau)} = \kappa \int \sum_{J}^* d\tau' \ p_{\tau}(\tau') \ \tanh^2(\beta J \tau \tau') \ \overline{h^2(\tau')} - \frac{\kappa - 1}{2} \overline{h(\tau)}^2$$
 (2.42)

For $\overline{h(\tau)}=0$ we are interested in the second equation, that defines a linear integral operator $f(\tau)=\int d\tau' \; \Gamma(\tau',\tau)f(\tau')$, with $f(\tau)=\overline{h^2(\tau)}$ and the (non-symmetric) kernel

$$\Gamma(\tau',\tau) = k \sum_{J=\pm J}^{*} p_{\tau}(\tau') \tanh^{2}(\beta J \tau \tau') . \qquad (2.43)$$

Therefore a solution of Eq. (2.42) with a non-vanishing function $\overline{h^2(\tau)}$ only exists if such integral operator has an eigenvector with eigenvalue 1. For the specific case of the box distribution of width Δ we have diagonalized the integral operator numerically for J=1 and $\kappa=2$, for several values of β and Δ , on a grid of 2048×2048 intervals. The results for the critical temperature are reported in blue Fig. 2.20. Note that for $\Delta\to 0$ the critical temperature tends to the Bethe lattice value for the Ising spin glass, $\beta_{\rm SG}J={\rm atanh}(\sqrt{1/\kappa})\simeq 0.88$, that is $T_{\rm SG}/J\simeq 1.1346$. The increasing trend is the same as the one derived in the fully connected model with the TAP method. The range of variation of the critical temperature with Δ is very weak in both cases. This growing dependence also holds for the clean version of the model with $J_{ij}=J>0$ in mean-field.

Although the behavior of the critical temperature differs between mean-field and finite-dimensional systems, the spin glass phase remains present in the former, as Δ is bounded. Combined with the frustration introduced by the symmetric bond distribution, this suggests that the spin glass order observed in d=2 persists in this model at T=0. The precise nature of these disordered bonds will be analyzed in later sections.

2.4.3 Dynamical properties toward the ground state

In this section, we will study the dynamical properties of the frustrated Δ -model. Specifically, we will characterize the evolution of a infinite temperature configuration—where both σ_i and τ_i are drawn randomly from their respective probability distributions—toward the equilibrium state at T=0, where spin glass ground states are present.

In this section, to different dynamical evolutions are considered:

- (i) An instantaneous temperature quench. Where the initial configuration is evolved with the Metropolis' acceptance rule at T=0.
- (ii) A temperature annealing, in which we slowly decrease the temperature according to a time-dependent protocol until reaching the target temperature [213], here being T=0.

The time-dependent protocol adopted in this work is given by

$$T(t) = T_0 \left(1 - t/t_f \right)^a \,, \tag{2.44}$$

with $T_0 = 1.0$, t_f the total number of MC-sweeps, and a = 1 (linear) or a = 2 (quadratic). Unless otherwise stated, we will set a = 2 as it seems to be more efficient than a = 1 for the case considered.

The two-time correlation

First, we have calculated the two-time correlation function of the soft-spin variables $s_i(t)$, defined as

$$C(t, t_w) = \frac{1}{N} \sum_{i=1}^{N} \left[\langle s_i(t) s_i(t + t_w) \rangle \right]$$
 (2.45)

where $s_i(t) = \sigma_i(t)\tau_i(t)$ for the Δ -model and $s_i(t) = \sigma_i(t)$ for the EA one. As can be seen in Fig. 2.21 the SWAP method induces a faster decay, compared to the single spin flip dynamics, in which the curves saturate rather quickly, consistent with plateau the R_p found in Fig. 2.2.

However, under the equivalent EA model interpretation, the diffusion of the τ_i variables in the SWAP evolution results in the partial annealing of the bonds \mathcal{J}_{ij} . Consequently, it is particularly illuminating to isolate this effects and understand the time evolution of the bonds \mathcal{J}_{ij} —through the spin length variables τ_i .

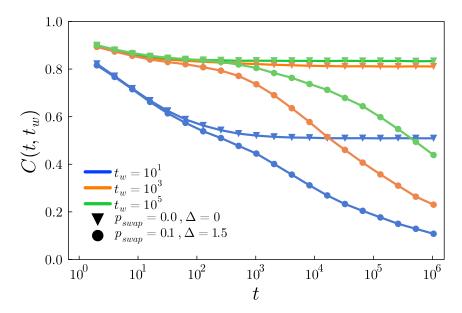


Fig. 2.21: The two-time correlation of the s_i variables, confronting the SWAP method ($p_{\text{swap}} = 0.1$) with the $\Delta = 1.5$ -model with the single-spin-flip kinetics of the 2d EA model, both with sizes L = 512 and the waiting time values displayed in the key.

Evolution of the effective bonds \mathcal{J}_{ij}

Frustrated plaquettes remain frustrated since neither the signs of J_{ij} nor \mathcal{J}_{ij} change. However, whenever SWAP is used the diffusion of the τ_i can affect the magnitude—or strength—of the local frustration, quantified by

$$f_P \equiv \prod_{\langle ij \rangle \in P} \mathcal{J}_{ij}(t) = \prod_{\langle ij \rangle \in P} J_{ij} \tau_i(t) \tau_j(t) , \qquad (2.46)$$

for frustrated plaquettes (i.e. $f_P < 0$). The product runs over the links of forming the plaquette P. Concretely, on a square plaquette with site labels 1, 2, 3, 4,

$$f_P(t) = J_{12}J_{23}J_{34}J_{41}\tau_1^2(t)\tau_2^2(t)\tau_3^2(t)\tau_4^2(t).$$
(2.47)

The sign, and whether the plaquette is frustrated or not, is decided by the factor $J_{12}J_{23}J_{34}J_{41}$ while the magnitude of the potential frustration is determined by $\tau_1^2(t)\tau_2^2(t)\tau_3^2(t)\tau_4^2(t)$ which depends on time.

The cumulative probability, defined as

$$\mathbb{P}^{<}(f_{P} < x) = \int_{-\infty}^{x} p_{f_{P}}(y) dy , \qquad (2.48)$$

is plotted for negative f_P at different three times reached with SWAP dynamics after a T=0 quench in the inset of Fig. 2.22. SWAP reduces the magnitude of the frustration, as the probability of finding large negative values for f_P decreases with time, up until a constant functional form is reached. Moreover, the two-time local correlation,

$$C_{\mathcal{J}}(t, t_w) = \frac{\sum_{i,j} [\langle \mathcal{J}_{ij}(t) \mathcal{J}_{ij}(t_w) \rangle]}{\sum_{i,j} [\langle \mathcal{J}_{ij}^2(t_w) \rangle]}, \qquad (2.49)$$

displayed in the main part of Fig. 2.22, becomes stationary and $\lim_{t\gg 1}\lim_{t_w\gg 1}C_{\mathcal{J}}(t,t_w)=1$. This can also be confirmed by tracking the energy density as a function of time, where a stationary plateau is reached at the same values of $t-t_w$. Effectively quenched configurations of the effective couplings $\mathcal{J}_{ij}^*=\mathcal{J}_{ij}(t_{\max})$ are reached in each run after $\sim 10^5$ sweeps in a system with L=32.

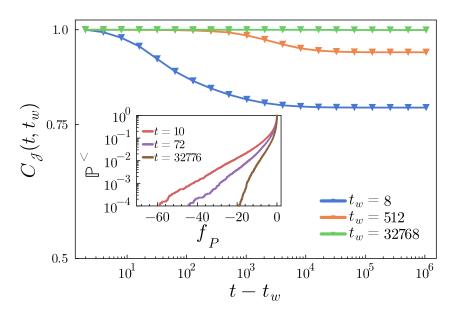


Fig. 2.22: Two-time correlations of the bonds \mathcal{J}_{ij} in the Δ -model with $\Delta=1.5$, L=32 quenched to T=0 evolved with SWAP ($p_{\mathrm{swap}}=0.1$). The waiting times t_w are given in the key. Inset: the cumulative probability of local frustrations f_P at three times after the quench.

The main conclusion of Fig. 2.22 is that for a zero temperature quench, after a time t_{\max} has elapsed the dynamics freezes, and the effective bonds $\mathcal{J}_{ij}(t_{\max}) = \mathcal{J}_{ij}^*$ are quenched. A natural question to ask is: what happens to the Ising spin configurations $\{\sigma_i\}$ of the equivalent EA interpretation? Are they ground states of the associated \mathcal{J}_{ij}^* ?

The ground states of \mathcal{J}_{ij}^*

We investigate the efficiency of SWAP to reach these ground states of the 2d EA model with the interactions \mathcal{J}_{ij}^* obtained at the latest time reached after a zero temperature quench with SWAP—where they have become stationary.

With this aim we stored the couplings $\mathcal{J}_{ij}(t)$ and the Ising spins $\sigma_i(t)$. Concomitantly, we used the facility in Bonn [234] to find the unique (apart from global spin reversal) ground state σ_i^{gs} of a 2DEA model with the \mathcal{J}_{ij}^* interactions—the functioning of this algorithm is

outlined in App. A1. Then, we calculated the overlap and the probability of reaching the ground state [235] as

$$q(t) = \frac{1}{N} \sum_{i=1}^{N} \sigma_i^{\text{gs}} \sigma_i(t) , \quad \mathbb{P}_0(t) = \frac{1}{N_r} \sum_{\alpha=1}^{N_r} \delta_{|q_\alpha(t)|,1} . \tag{2.50}$$

The index α runs over the simulation runs and N_r is its total number, that is, the total number of MC simulations which sample different realizations of the couplings J_{ij} , the set of lengths $\{\tau_i\}$ and the initial conditions $\{\sigma_i(t=0)\}$. In the simulations shown this number was always around 100. The results are shown in Fig. 2.23 where we have denoted the plateau reached at large t as \mathbb{P}_0^{∞} , meaning the asymptotic probability of finding a ground state.

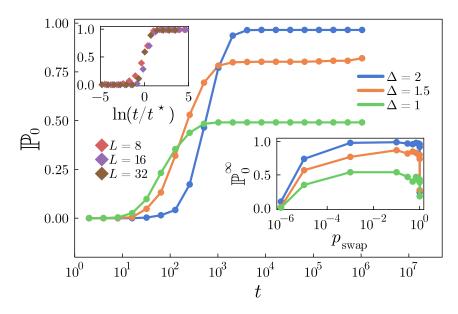


Fig. 2.23: The probability of reaching the ground state of the Δ -model with $\mathcal{J}_{ij}^* = \mathcal{J}_{ij}(t_{\rm max})$, in systems with L=8 after zero temperature quenches. Main panel: three Δ values given in the right key and $p_{\rm swap}=0.5$. For $\Delta=2$, 96.5% of the runs find a ground state. Lower inset: the asymptotic value against the parameter $p_{\rm swap}$ for $\Delta=2$. Upper inset: Scaling with $t^*(L)=1.25\,L^{3.75}$ of the data for $\Delta=2$, $p_{\rm swap}=0.1$ and three system sizes specified in the left key.

For an instantaneous quench to T=0, the main panel in Fig. 2.23 displays $\mathbb{P}_0(t)$ for three values of Δ and $p_{\mathrm{swap}}=0.5$. In all cases, after a fast increase ending at $t\lesssim 10^3$ MCs, $\mathbb{P}_0(t)$ saturates to a \mathbb{P}_0^∞ which increases with Δ and gets very close to 1 for $\Delta=2$. Further optimization of the algorithm achieved by gauging p_{swap} is studied in the lower inset which shows the dependence of \mathbb{P}_0^∞ on p_{swap} in the Δ -model with $\Delta=2$. For intermediate values, $0.1\lesssim p_{\mathrm{swap}}\lesssim 0.9$, \mathbb{P}_0^∞ remains approximately constant apart from numerical noise, and it decays to zero at the two extremes of either non-local spin exchanges $(p_{\mathrm{swap}}\to 1)$ or pure single-spin-flips $(p_{\mathrm{swap}}\to 0)$, as a result we will use $p_{\mathrm{swap}}=0.1$ for most subsequent analysis. Finally, we checked the dependence on system size in the lower panel using $\Delta=2$ and $p_{\mathrm{swap}}=0.1$. The curves are similar and the percentage of ground states found is independent of L. The upper inset shows the scaling of \mathbb{P}_0 against $t/t^*(L)$ with $t^*(L)\sim 1.25\,L^{3.75}$.

An inconvenience resides in the fact that even in the cases in which the ground states are reached, for the same realization of the J_{ij} and initial conditions of the τ_i , different thermal evolutions produce different final configurations with a broad energy density distribution, as can be seen in the inset in Fig. 2.24(b) below. Therefore, although being ground states of the 2DEA model with couplings \mathcal{J}_{ij}^* , these configurations originate from metastable states of

the τ_i variables. We are interested in finding a way of closing this gap, and thus finding a global minimum of the full model in Eq. (2.28).

In order to optimize both the τ_i and σ_i , we adopt the thermal annealing protocol defined in Eq. (2.44)

$$T(t) = T_0 \left(1 - t/t_f \right)^a , \qquad (2.51)$$

with $T_0 = 1.0$, t_f the total number of MC-sweeps, and a = 2.

SWAP is now able to find ground states 100% of the runs for almost all Δ values, as seen in the pink triangles in Fig. 2.24(a). The spread of ground state energies (Gaussian distributed) narrows considerably with respect to the one of T=0 quenches, in inset Fig. 2.24(b) (although it does not disappear completely). The configurations reached are ground states of models with only slightly different \mathcal{J}_{ij}^* . This is confirmed by the evolution of the self correlation of the length variables in two different runs, $\tau_i^{(1)}$ and $\tau_i^{(2)}$ that will be shown below in Fig. 2.25.

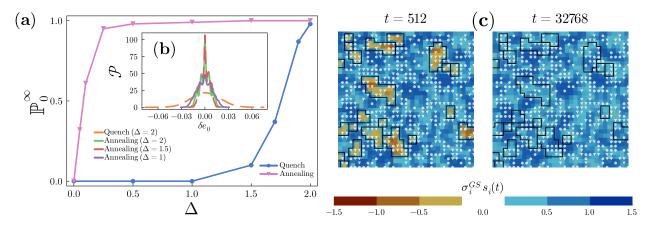


Fig. 2.24: (a) Asymptotic probability of reaching a ground state after a T=0 quench and a quadratic annealing, starting from $T_0=1.0$ during $t_f\approx 10^7$ MCs. L=32. Inset (b) Probability distributions of the ground state energy density differences found after T=0 quenches and annealing protocols, Eq. (2.44), of models with different Δ . In all cases, the J_{ij} and initial lengths $\{\tau_i(t=0)\}$ are the same, and the data are sampled over 10^3 initial Ising spin conditions $\{\sigma_i(t=0)=\pm 1\}$. (c) The overlap of an early (left panel) and final (right panel) s_i configuration with the ground state of the model with couplings $\mathcal{J}_{ij}^*=\mathcal{J}_{ij}(t_{\rm max})$, for a quench to T=0 with $\Delta=1.5$. The light bullets and triangles are located at frustrated plaquettes with local frustration f_P being greater or smaller than one-half in magnitude, respectively.

The snapshots shown in Fig. 2.24(c) confirm the initial hypothesis illustrated in Fig. 1 and reveal more detailed mechanisms. By analyzing the overlap of an evolving configuration with its respective expected ground state, a quench at T=0 displays clear coarsening behavior. Warm colors (enclosed with black lines) denote droplets, regions where spins are unaligned with the expected ground state, while the cool colors indicate domains where spins are properly aligned. We label the frustrated plaquettes ($f_P < 0$) at their centers with circles (for $|f_P| > 0.5$) and triangles (for $|f_P| \le 0.5$) depending on the magnitude of the frustration strength. At t=512, the frustration decreases in several of the enclosed droplets (more triangles accumulate around the droplets), favoring spin flips that will destroy the droplets and align their spins with the ground-state orientations. After this process, when the ground state has been reached at t=32768, the newly ground-state-aligned spins can acquire a larger magnitude (lowering the energy further) explaining the proliferation of circles in the regions previously occupied by the droplets.

Evolution of the τ_i variables

We show here how the temperature annealing recovers similar \mathcal{J}_{ij}^* by analyzing the behavior of the τ_i variables. Specifically, we track the correlation of the τ_i variables sampled in different runs of the dynamics, given by

$$C_{\tau^{(1)}\tau^{(2)}}(t) = \frac{\left[\frac{1}{N}\sum_{i=1}^{N}\tau_{i}^{(1)}(t)\tau_{i}^{(2)}(t)\right] - 1}{\left(1 + \frac{\Delta^{2}}{12}\right) - 1} . \tag{2.52}$$

In Fig. 2.25 we plot the time evolution of this correlation, where $\tau_i^{(1)}$ and $\tau_i^{(2)}$ are the values of the length variables in two runs of the model with the same quenched disorder J_{ij} , starting from the same initial condition of the Ising spins and length variables. At equal times $C_{\tau^{(1)}\tau^{(2)}}(t=0)=1$ by definition since $\tau_i^{(1)}(0)=\tau_i^{(2)}(0)$ and $[\tau_i^2]=1+\Delta^2/12$. If for $t\to\infty$ the $\tau_i^{(1)}$ and $\tau_i^{(2)}$ lengths fully correlate again, then 1 should be recovered in this limit as well. This is confirmed in Fig. 2.25 where this time-dependent correlation is shown for various values of Δ in a system with L=16 annealed following the protocol Eq. (2.51) from $T_0=0.5$ to zero temperature.

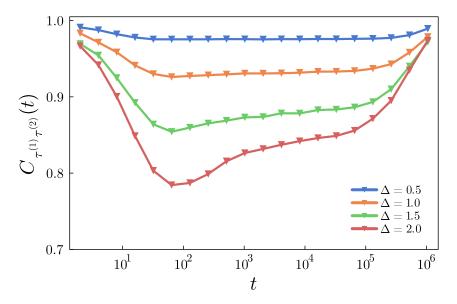


Fig. 2.25: The correlation of the τ_i variables according to Eq. (2.52). L=16 system annealed to zero temperature.

The evolution of the frustration magnitude

The mechanism that renders the SWAP method efficient in the frustrated Δ -Model is clarified by inspecting some snapshots, as we did for the ferromagnetic model. In Fig. 2.27, we show the actual configuration $\{s_i\}$ (first row) and the overlap with the ground state of the system with the final \mathcal{J}_{ij}^* interactions $\{\sigma_i^{\mathrm{gs}}s_i\}$ (second row). We have performed a sub-critical quench at T=0, for L=128 using SWAP dynamics with $p_{\mathrm{swap}}=0.1$.

The first row displays the spin configurations at four times after the quench. The images do not show much structure apart from a slight tendency of long spins with the same orientation to group locally. Still, and as expected, no structure with long-range spin ordering develops.

One can recognize the formation of domains in the overlap of the configurations with the ground state. Indeed, the images shown in the second row are primarily red or yellow, that

is, the system acquires a positive overlap with the ground state all over space. The SWAP algorithm produces domain walls consisting mostly of short length spins (green and yellow). This structure lowers the energy barriers locally, making the spin-flip more feasible in these regions. The ratio of $\{\tau_i\}$ variables for which the corresponding Ising spins do not match the ground state (i.e. $\sigma_i(t)\sigma_i^{gs}<0$) goes to zero, see Fig. 2.26, in the course of time. However, the decay rate depends strongly on the length of the spins. Larger spins $(1<\tau_i\leq 2)$ align faster on the ground state directions compared to smaller ones $(\tau_i\leq 1)$, as can also be seen in Fig. 2.26, supporting the previous explanation.

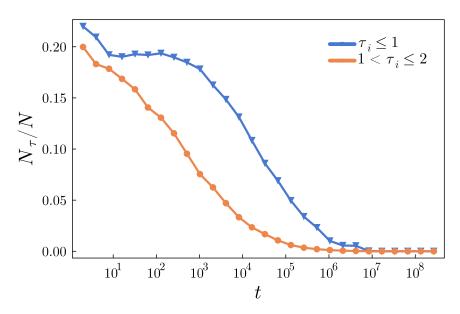


Fig. 2.26: Time evolution for the ratio of small-length ($\tau_i \leq 1$) and large-length ($1 < \tau_i \leq 2$) spins with σ_i not aligned with the ground state $\sigma_i^{\rm gs}$, i.e. $\sigma_i(t)\sigma_i^{\rm gs} < 0$.

From the snapshot in Fig. 2.24(c)—and several others shown in App. A2—it is clear that the frustrated or unfrustrated nature of the plaquettes does not change in time. When a symbol, such as a bullet or triangle, is attached to them, they neither disappear nor appear elsewhere.

However, the strength of local frustration varies over time. As a result, the symbols may change—for example, from a bullet to a triangle, or vice versa. This behavior is characterized by the number density of frustrated plaquettes, n, whose time evolution is shown in Fig. 2.28. We classify the plaquettes into two categories: those with large amplitudes ($|f_P| > 1$) and those with small amplitudes ($|f_P| < 1$). The number of plaquettes in each category is normalized by the total number of frustrated plaquettes, N_F . Notably, N_F remains constant over time and is approximately equal to half of the total number of plaquettes.

The number density of plaquettes with small frustration dominates over the ones with large frustration asymptotically, making the magnitude of the total frustration to diminish in time. This is better quantified by measuring the total frustration density f(t), defined by

$$f(t) = \frac{1}{N_F} \sum_{P} f_P(t) , \qquad (2.53)$$

with the sum running over frustrated plaquettes.

The time evolution of this quantity is shown in Fig. 2.29 for different values of Δ where, as shown in the inset of Fig. 2.22, the overall frustration strength diminishes. For $\Delta=2$ —at which the SWAP method with a quench is most efficient—the asymptotic value of f(t) saturates at a value $\gtrsim 1$.

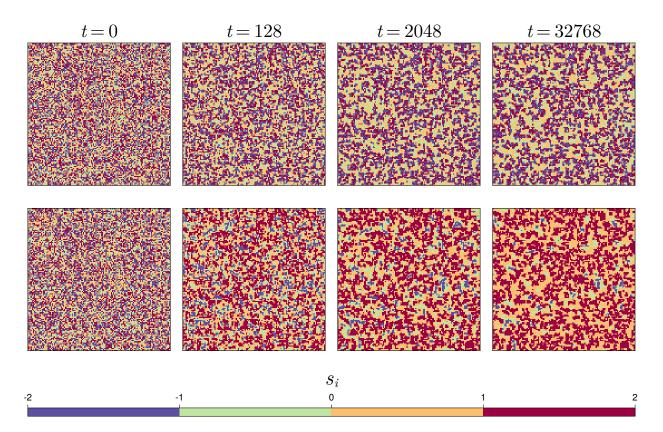


Fig. 2.27: Instantaneous snapshots of the frustrated Δ -model with $\Delta=2$, at four times indicated in the figure after a quench to T=0. First row: The spin s_i configurations. Second row: Overlap with the σ -ground state $(s_i(t)\sigma_i^{\rm gs})$. The color scale binned as indicated in the bar shows the lengths of the local spins.

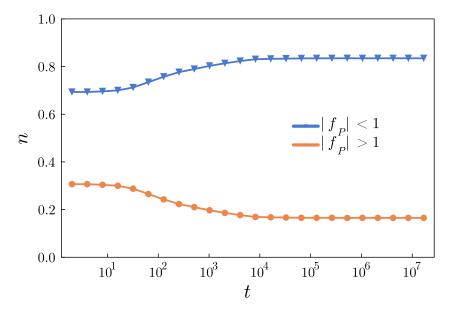


Fig. 2.28: The number density of plaquettes with frustration $f_P < 0$, distinguished by the modulus being larger or smaller than 1.

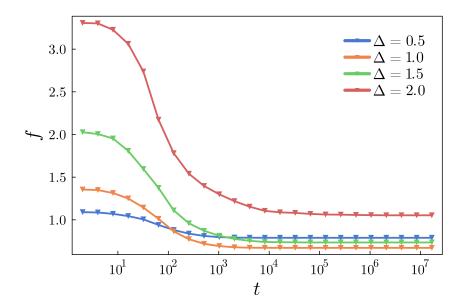


Fig. 2.29: The magnitude of the total frustration f, defined in Eq. (2.53) in Δ -models with different values of Δ specified in the key and L=32 quenched to zero temperature and evolved with SWAP. The initial value is $(1+\Delta^2/12)^4$.

The stationary bond distribution

Overall, these subtleties lead to the converged bonds \mathcal{J}_{ij}^* . In Fig. 2.19 we compare the bond distribution at an early stage (t=2), which closely matches the theoretical prediction, with the distribution at a much later stage $(t=2^{24})$, where the bonds have already reached a stationary state. These results were obtained from averages over $N_r=100$ independent runs following a quench to after a quench at T=0.

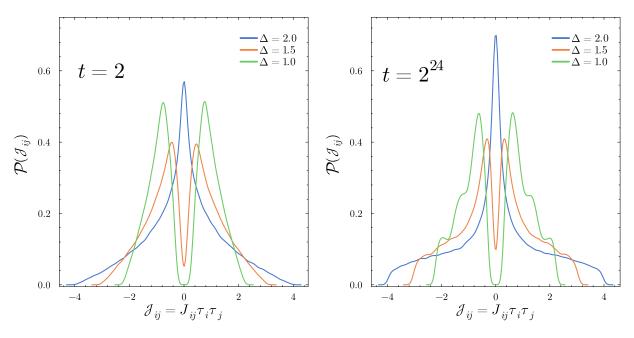


Fig. 2.30: Distribution of the effective bonds \mathcal{J}_{ij} at two instants of the simulation: (almost) initial distribution (t=2), final distribution ($t=2^{24}$) after a SWAP method was performed ($p_{\text{swap}}=0.1$). Taken from an evolution at T=0 quench, with L=32, and the values of Δ displayed in the key.

The distribution preserves its symmetric structure. At later times the tails widen for all

 Δ and symmetric peaks arise, always lower than the original typical values, that are still present. No new peaks arise for $\Delta=2$, however the original peak increase its probability, reflecting the overall reduction of frustration strength.

2.4.4 Dynamical properties at finite temperature

We now examine the properties of the algorithm at finite temperature. To do this, we measure the overlap—overlap correlation function, which allows us to extract the growing length R(t). As discussed in Sec. 2.2, this length saturates $T \leq 0.2$, and the dynamics freezes as illustrated in Fig. 2.2.

For clarity, we recall the definition of the overlap—overlap (also referred to as the four-point spin—spin) correlation function, and the extracted spin glass growing length:

$$C_4(r,t) = \frac{1}{N} \sum_{\substack{i,j=1\\|\vec{r}_i - \vec{r}_j| = r}}^{N} \left[\left\langle \sigma_j^{(1)}(t) \sigma_j^{(2)}(t) \sigma_i^{(1)}(t) \sigma_i^{(2)}(t) \right\rangle \right], \ R_{\sigma}(t) = 2 \int_0^{\infty} dr \ C_4(r,t) \ , \quad (2.54)$$

The spatial dependence of C_4 , at different times, is studied in Fig. 2.31 at T=0.5 and T=0.1. At high temperatures SWAP on the $\Delta=1$ model yields equivalent results to the single-spin flip evolution of the original $\Delta=0$ model. At the largest time $t\sim 10^6$ the growing length is of the order of 20, say, while the system's linear size is L=512.

At low temperatures instead, SWAP is much more efficient in building long-range correlations, even between copies that may have evolved toward different effective random couplings \mathcal{J}_{ij} , than the single-spin flip evolution of the original $\Delta=0$ model. Still, the correlations obtained with SWAP at the same time $t\sim 10^6$ decay faster than at higher T, reaching a shorter $R_{\sigma}(t)$.

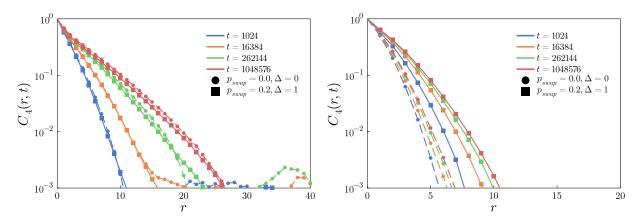


Fig. 2.31: Overlap correlation $C_4(r,t)$ as a function of the Cartesian distance r for several times given in the keys, with $L=512\pm J$. 2d EA ($\Delta=0$) model evolved with single-spin dynamics (circles) and $\Delta=1$ model evolved with SWAP and $p_{\rm swap}=0.2$ (squares). Data for T=0.5 (left) and T=0.1 (right). While single spin flips and SWAP yield equivalent results at high temperature, the latter builds longer correlations at low temperature.

Finally, in Fig. 2.32 we compare the performance of SWAP with different values of the parameter $p_{\rm swap}$ to the one of the single spin flip updates when applied to the $\Delta=1$ model at high (left) and low (right) temperatures. The data in the figure demonstrate that the two methods perform similarly for high temperatures.

However, for low temperatures (T=0.1) the single spin flip dynamics of the $\Delta=0$ model freeze and the spin glass length saturates to $R_{\sigma}\sim 5$ (blue data). Instead, both single spin flips (orange) and SWAP (green) of the $\Delta=1$ model accelerate the dynamics at long

times, with the latter becoming more efficient in the last four time decades. We estimated running effective exponents from fits over a moving window with 12 data points (inset). At such low temperature within a numerical accuracy $z_{\rm eff}$ converges to $z_{\rm eff}^{\infty} \sim 49$. The dashed line is the algebraic law $A\,t^{1/z_{\rm eff}^{\infty}}$. Long-time configurations are not in equilibrium. First, the

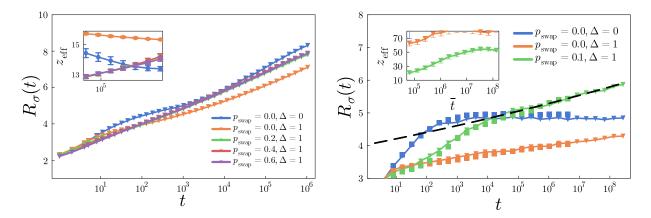


Fig. 2.32: Typical growing length of the frustrated model with SWAP dynamics, the values of Δ and $p_{\rm swap}$ are shown in the keys. We show the evolution at (left) a relatively high temperature T=0.5 with L=512 and (right) at a lower temperature T=0.1 using both L=32 (triangles) and L=512 (squares). The growth in the $\pm J$ 2d EA with single spin flip dynamics is plotted for comparison (blue curves). The insets show the time evolution of the dynamical exponent. In dashed black is the fit with $z_{\rm eff}=49.5(86)$ and amplitude A=3.94(2)

maximal length $R_{\sigma} \sim 6$ for L=32 is far from L/2. Second, the τ_i variables, which are also dynamical under SWAP, are still evolving. Hence, $\sigma_i^{(1)}$ and $\sigma_i^{(2)}$ may be optimized with respect to different \mathcal{J}_{ij} and thus not really inform us about the performance of the algorithm and measurement in taking one system close to equilibrium.

2.4.5 Consolidating the key results

Our SWAP has two ingredients: annealing of disorder because of the τ_i exchanges (absent in the original SWAP method used for interacting particle systems) and non-local moves (the essential feature of the original SWAP). To determine whether the annealing of the τ_i —and the resulting reduction in the magnitude of frustration—yields disorder realizations \mathcal{J}_{ij}^* with a less complex structure, in which the ground state can be found more easily, we compared the outcome of (i) our SWAP algorithm to the following cases:

- (ii) Purely single spin flip dynamics (i.e. $p_{\text{swap}} = 0$) with randomly generated \mathcal{J}_{ij} bonds. They follow the initial distribution of Fig. 2.19, with uncorrelated $\{\tau_i\}$.
- (iii) Purely single spin flip dynamics with the converged \mathcal{J}_{ij}^* , found with the quadratic temperature annealing using the strictly non-local SWAP method (with $p_{\text{swap}} = 0.1$). In this case the τ_i variables are quenched but organized in the pattern produced by the SWAP implementation.

We also wanted to distinguish the effect of the non-local moves, therefore we included:

(iv) An evolution with only local exchanges, in which we restrict the spin exchanges to be just between nearest neighbors. As with the non-local SWAP implementation the τ_i 's are still partially annealed along the evolution but only through local moves.

We then calculate the self-correlation $C_{\sigma}(t,t_w)=N^{-1}\sum_{i=1}^N[\langle\sigma_i(t)\sigma_i(t_w)\rangle]$ after a T=0 quench of the L=32 system at three waiting times. The results are plotted in Fig. 2.33, they show that single spin flips, for both \mathcal{J}_{ij} and \mathcal{J}_{ij}^* , fail to decorrelate the configurations, as their curves lie in the plateau of $C(t,t_w)=1$ for all waiting times. The \mathcal{J}_{ij}^* are not special in this respect. Besides, while local exchanges are able to decorrelate configurations, they also reach the $C(t,t_w)=1$ plateau at the last waiting time considered ($t_w=512$), but the associated configuration of the σ_i variables is not the ground state of the converged \mathcal{J}_{ij}^* , unlike for the non-local SWAP. Thus, non-local SWAP is more efficient in advancing the evolution than just local exchanges.

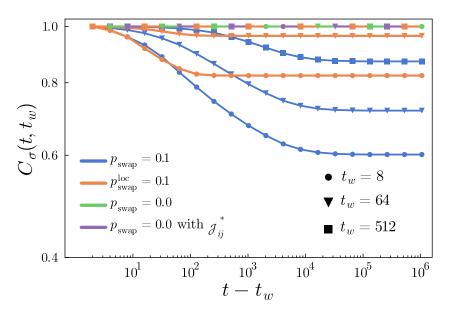


Fig. 2.33: The two-time Ising spin self-correlation at three waiting-times. Data for four kinds of T=0 evolution of a system with L=32 and $\Delta=1.5$ starting from random initial conditions: (i) SWAP with non-local moves, (iv) local spin exchanges and spin flips, and solely single spin flips of the Δ model with (ii) random \mathcal{J}_{ij} and (iii) converged \mathcal{J}_{ij}^* couplings.

Finally, we measured the characteristic relaxation time, τ_{α} , of the algorithms at finite temperature. We have excluded algorithm (ii), the case of pure single spin flips with randomly generated \mathcal{J}_{ij} bonds, because our previous analysis concluded that its dynamics are similar to those with converged bonds, \mathcal{J}_{ij}^* . This time the converged couplings are taken from the values produced at the end of the non-local SWAP quench when equilibrium is reached, for each respective temperature. We define this τ_{α} as the time for which the self-correlation has become age independent (i.e. $C(t,t_w)=\hat{C}(t-t_w)$) and has decayed to 20% of its original value (i.e. $\hat{C}(\tau_{\alpha})=0.2~\hat{C}(0)$). As can be seen in Fig. 2.34, the partial annealing of the disorder plays a prominent role in accelerating the dynamics with respect to the single spin flip case with converged couplings. The addition of non-local moves accelerates the dynamics even further, gaining one decade at T=0.9 with respect to the local exchanges. The three dynamics are indistinguishable at higher temperatures (here, around $T\approx 2.5$).

2.5 Conclusions and outlook

In this work, we extend standard lattice spin models by introducing associated Δ -models, in which spins are assigned a length variable drawn initially from a predetermined probability distribution with positive support (here chosen to be uniform). Within this framework, we

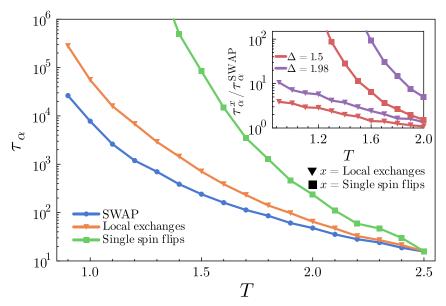


Fig. 2.34: Characteristic relaxation time τ_{α} extracted from the decay of the σ self-correlation after quenches to the target temperatures, evolved with (i) SWAP (i.e. non-local moves), (iv) local spin exchanges, both with the same $p_{\rm swap}=0.1$, and (iii) pure single-spin-flip dynamics ($p_{\rm swap}=0$) with converged bonds (\mathcal{J}_{ij}^*), for a model with $\Delta=1.5$. In the inset, comparison of the two relaxation times, for local exchanges and pure single-spin-flips with respect to SWAP, as a function of temperature for two Δ -models.

adapted the SWAP method to perform non-local spin exchanges, aiming to evaluate how these moves help overcome local energy barriers in the Metropolis algorithm.

Our implementation is based on single-spin-flip dynamics, which constitute the majority of updates. In addition, non-local spin-exchange (SWAP) moves are introduced at random intervals, accounting for roughly 10% of the total updates—when optimized. This procedure preserves the overall distribution of spin magnitudes but induces emergent spatially correlated patterns. These correlations effectively lift local energy barriers, thereby facilitating the acceptance of new configurations.

In the Δ -model with uniform ferromagnetic couplings, the SWAP mechanism offers no advantage over the standard single-spin-flip dynamics of the parent Ising model. Instead, it merely reproduces the well-known non-conserved order-parameter kinetics of the Ising model, effectively bypassing the disorder introduced by the τ_i variables.

In contrast, in the frustrated case, SWAP manages to accelerate the evolution of a 2d frustrated model at very low temperatures. The method allowed us to sample ground states of the 2d Ising spin glass, faster than usual implementations of parallel tempering do for the EA model at low temperatures. After instantaneous quenches, we find the ground state of the Δ -Model with $\Delta=2$ around 2 decades faster than parallel tempering Monte Carlo [235, 236] in the 2d EA model. This is further improved by a temperature annealing, that enables us to find at least 99% of the ground states for $\Delta \geq 0.5$. Other parameters to optimize, which we have not explored in depth yet, are $P_{\tau}(\tau_i)$ and the parameters of the annealing scheme.

In the disordered spin model, the sluggish dynamics at low temperatures stems from the complex organization of low-energy configurations and the difficulty of transitioning between them. The efficiency of the SWAP algorithm in accelerating the dynamics can therefore be attributed to its ability to partially reduce these transition difficulties, effectively mitigating the largest energy barriers when needed. This interpretation is supported by the absence of any speed-up in the non-disordered Ising case, whose low-energy configurations are comparatively simple. Nevertheless, the acceleration achieved in our spin model is

modest compared to that observed in structural glasses, suggesting that facilitation [184] may also contribute in the latter, consistent with previous studies of kinetically constrained models [77]. In addition, we observe correlated patterns of spin lengths, τ_i , at low temperatures, hinting at similar phenomena in structural glasses, where spatial correlations related to particle diameters may develop—an effect that merits further investigation.

A future direction is to extend the application of this algorithm to the 3d case or random regular graphs (RRG), where the spin glass transition occurs at finite temperatures. The RRG case is particularly interesting, as we expect the critical temperature of the Δ -model to be well estimated by the calculation on the Bethe lattice. This will allow us to assess the dependency $T_c(\Delta)$ in finite-dimensional systems.

Although the frustration magnitude is reduced overall, the probability distribution of the converged couplings \mathcal{J}_{ij}^* remains symmetric and does not appear to constitute an 'easier instance' than standard alternatives such as Gaussian or bimodal distributions. A more complete characterization of these random bonds would require estimating the associated stiffness exponent, known to be $\theta = -0.2793(3)$ for continuous distributions [237, 238]. In this work, we could not estimate this exponent directly, as the ground states generated by SWAP correspond to different \mathcal{J}_{ij}^* instances, preventing the standard domain-wall analysis of ground state energies. However, it could still be determined by computing ground states of the \mathcal{J}_{ij}^* instances produced by SWAP using conventional methods.

A1 Branch-and-cut algorithm

Currently, the branch-and-cut algorithm can be considered as the most efficient exact algorithm to obtain ground states for spin glass instances. A publicly accessible implementation of a branch-and-cut algorithm is hosted by the University of Bonn [234]. In this work, we have used this server to validate the ground state sample produced by the SWAP algorithm.

The algorithm is employed to find the global energy minima of a general Ising spin glass with Hamiltonian

$$\mathcal{H} = -\sum_{i < j} J_{ij} \sigma_i \sigma_j - \sum_i h_i \sigma_i , \qquad (55)$$

with J_{ij} being quenched random variables—either continuous or discrete.

For system of size N, the naive algorithm would compute the energy of the 2^N possible spin configurations and select the minimum from this list, producing a time complexity of $O(2^N)$. Instead, we follow the branch-and-cut construction, that simplifies the problem to a recursion.

Depending on the value of the spin σ_N the Hamiltonian can be written in two ways:

$$\mathcal{H}_{+}(\boldsymbol{\sigma}') = -\sum_{i < j}' J_{ij}\sigma_{i}\sigma_{j} - \sum_{i}' h_{i}\sigma_{i} - \sum_{i}' J_{iN}\sigma_{i} - h_{N}, \qquad (56)$$

$$\mathcal{H}_{-}(\boldsymbol{\sigma}') = -\sum_{i < j}' J_{ij}\sigma_{i}\sigma_{j} - \sum_{i}' h_{i}\sigma_{i} + \sum_{i}' J_{iN}\sigma_{i} + h_{N}, \qquad (57)$$

where the sum \sum' runs from 1 to N-1. By defining the minimum of the Hamiltonian without the N-th spin, i.e.

$$\mathcal{H}'_{N-1} = \min_{\boldsymbol{\sigma}' = (\sigma_1, \dots, \sigma_{N-1})} \left[-\sum_{i < j}' J_{ij} \sigma_i \sigma_j \right] . \tag{58}$$

using the min-sum inequality $\min_x (f_1(x) + f_2(x)) \ge \min_x f_1(x) + \min_x f_2(x)$ we obtain the bounds

$$\mathcal{H}_{+}(\boldsymbol{\sigma}') \geq \mathcal{H}'_{N-1} - h_{N} + \min_{\boldsymbol{\sigma}' = (\sigma_{1}, \dots, \sigma_{N-1})} \left[\sum_{i}' (-h_{i} - J_{iN}) \sigma_{i} \right] ,$$

$$\mathcal{H}_{+}(\boldsymbol{\sigma}') \geq \mathcal{H}'_{N-1} + h_{N} + \min_{\boldsymbol{\sigma}' = (\sigma_{1}, \dots, \sigma_{N-1})} \left[\sum_{i}' (h_{i} - J_{iN}) \sigma_{i} \right] .$$
(59)

If we regard the original Hamiltonian $\mathcal{H}(\sigma)$ as the root of a graph, fixing one spin at a time produces a binary branching tree: the root splits into two vertices corresponding to $\sigma_N=+1$ and $\sigma_N=-1$; each of these branches again according to the possible values of σ_{N-1} , and so on. Continuing this process recursively generates a tree with 2^N leaves, where each leaf represents one possible spin configuration.

At each branching step, we obtain the lower bounds on the minimal energy attainable within each subtree found in (59). These bounds are then compared to a 'current best energy', which is obtained either from a heuristic algorithm or from previously explored branches. Any branch whose lower bound already exceeds this threshold is pruned, as it cannot contain the ground state. By iteratively branching, bounding, and pruning, the algorithm efficiently narrows the search and identifies the exact ground-state configuration at one of the surviving leaves.

In the worst case, the algorithm still scales exponentially, $O(2^N)$, since spin glasses are NP-hard. However, in practice, the effective search space is much smaller. Empirical studies

show that branch-and-cut can solve instances up to hundreds of spins efficiently, where naive enumeration would be impossible.

A2 More snapshots at L = 32

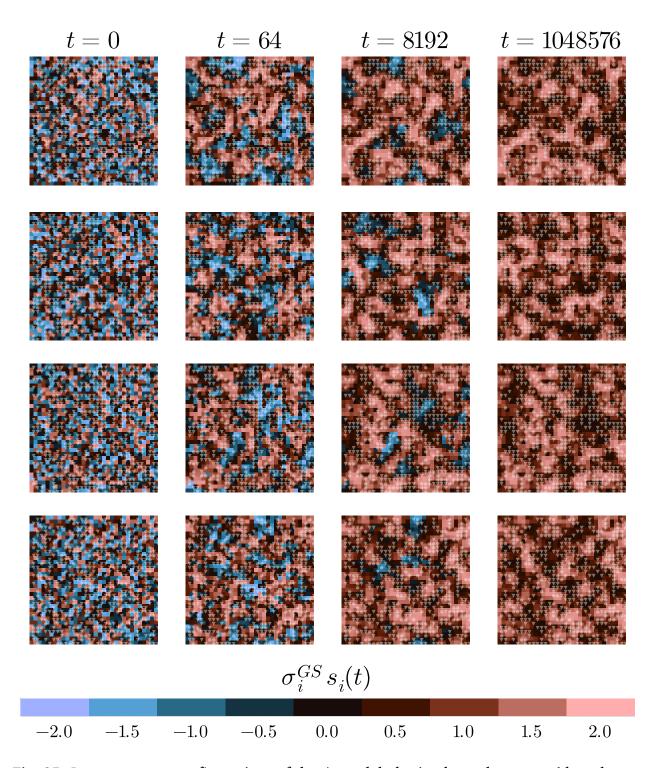


Fig. 35: Instantaneous configurations of the Δ -model obtained at subsequent (though not equally spaced) times after a quench to T=0 of four initial conditions (different rows) evolved with SWAP. Each square represents the instantaneous local overlap $\tau_i(t)\sigma_i(t)\sigma_i^{\rm gs}$. The frustrated plaquettes are indicated with bullets and triangles according to the local frustration $f_P(t)$ being greater or smaller than one, respectively.

Chapter 3

The Importance of Rare Events in Many-Body Localization

Contents

3.1			
	3.1.1	The Anderson model	112
	3.1.2	Eigenstate characterization	115
	3.1.3	Spectral Properties	116
		Relationship with the time evolution	
3.2			
	3.2.1	From Anderson to Many-Body Localization	122
	3.2.2	Perturbative Perspective on MBL	123
	3.2.3	Contrasting MBL with Other Forms of Ergodicty Breaking	124
	3.2.4		
	3.2.5	Quasilocal Integrals of Motion	126
3.3			129
	3.3.1	The Hilbert space picture	129
	3.3.2	Spin basis	131
	3.3.3		
3.4	Standard observables and their limitations		132
	3.4.1	The destabilization of the MBL phase	139
3.5	The large deviation method		142
	3.5.1	The delocalization probability	142
	3.5.2	A Proxy for Transport in Hilbert space	146
3.6	Finite size analysis		155
	3.6.1	The finite size phase diagrams	156
		Dependence on the target states distance	
	3.6.3	The distributions of the Hilbert space Landauer tranmssions	164
3.7	Rare vs typical samples		165
3.8	Concl	Conclusions and outlook	
A1	Detail	Details of the numerical simulations	
A2	The computation of the error bars		177
A3	The D	e DPRM correlations	
A4	Supplementary results for the Anderson basis		179
	A4.1	Dependence on the target basis states	179
	A4.2	The eigenstate amplitude decays	180

Introduction

Many-body localization (MBL) is a dynamical phase of matter in which an isolated, interacting quantum system with sufficiently strong disorder, controlled by a disorder strength W, fails to thermalize. Although the possibility of localization in interacting systems was first suggested by Anderson in his pioneering work [88], it was only about two decades ago that the perturbative stability of the Anderson insulator in the presence of weak interactions was explored [89, 90]. For recent reviews, see Refs. [91–94]. This failure is attributed to the emergence of an extensive set of quasi-local integrals of motion (LIOMs or ℓ -bits') that inhibit thermalization, allowing the system to retain local memory of its initial conditions indefinitely [91–94, 239–242].

In recent years, the stability (particularly with respect to *non-perturbative* events) of the MBL phase has been put into question [95–97, 100–103, 243–245]. This skepticism stems from the exponential increase in Hilbert space volume with the number of degrees of freedom, which severely limits exact diagonalization studies to relatively small system sizes, potentially missing the asymptotic behavior of large systems. In these studies, commonly used observables—such as average spectral statistics, eigenstate participation entropies, imbalance decay, and entanglement entropy [102, 246, 247]—fail to yield a consistent estimate of the critical disorder strength $W_{\rm MBL}$. As the system size L is increased, this putative critical point drifts toward larger disorder strengths, indicating that larger systems require stronger randomness to localize. This raised concerns that the true transition in the thermodynamic limit might occur only at infinitely strong disorder. In other words, the MBL regime seen in finite chains might be a finite-size effect rather than a stable phase [95–103].

This possibility has been further supported by a theoretical argument, which suggests that the MBL phase may be unstable with respect to a runaway avalanche thermalization mechanism. This instability would be triggered by rare regions in the system where the disorder is anomalously weak [100, 243, 248–254]. The basic idea is that rare regions with atypically weak disorder inevitably occur in sufficiently large systems, albeit with small but finite probability. In an otherwise localized system, rare regions can act as small thermal bubbles that couple to nearby degrees of freedom. Under certain conditions, a thermal bubble can thermalize its neighbors, which are then effectively absorbed into the bubble. As the bubble grows, it becomes increasingly capable of thermalizing additional nearby regions. This can trigger a self-sustained avalanche in which the thermal region expands throughout the system, ultimately destroying localization and restoring ergodic behavior.

Yet, the system sizes accessible through numerical simulations are too small to accommodate these rare, locally thermal regions, making direct evidence of their existence difficult to obtain. To address this limitation, several studies have investigated how disordered quantum chains in the MBL phase respond when coupled to an artificially introduced thermal region [248, 253–261]. The core idea is that if the relaxation time of the localized region is long enough, a thermal bubble cannot trigger the thermal avalanche. This approach sets a lower bound on the disorder strength required for the MBL phase to remain stable, and the critical disorder estimated turns out to be much stronger than those previously obtained using standard observables [244, 246, 247, 262, 263].

In parallel, many recent numerical studies of finite-size quantum systems have revealed the presence and importance of rare many-body resonances, even deep within regions of the phase diagram that were previously thought to belong to the many-body localized phase [264–268]. A many-body resonance occurs when two or more eigenstates of the system have almost the same energy and are connected strongly enough by the Hamiltonian that the system can transition between them. In the MBL phase, these eigenstates are typically 'localized' and differ only within a finite spatial region. When resonances involve eigenstates

that differ by a large fraction of the system's degrees of freedom (for example, several spin-flips in spin chain) the avalanche is said to be long-range (or system-wide) resonances, and are believed to be a key factor in its eventual breakdown. Evidence for such resonances has been found indirectly in microscopic models [268–272] and, more recently, identified directly through targeted numerical studies [257, 258, 265, 266, 273, 274] and play a central role in formal proofs about when the MBL phase can remain stable [264, 275, 276].

It is important to emphasize that these two destabilization mechanisms—many-body resonances and avalanches—are not necessarily independent. In particular, Ref. [258] argues that avalanches primarily propagate through strong, rare (near-)resonances, suggesting a close connection between the two. However, the precise relationship between avalanches and long-range resonances remains elusive. The growing body of evidence for many-body resonances places them at the center of the discussion, as they seem to play a crucial role in the physical properties of the MBL transition and its stability [251, 257, 258, 266, 270, 272, 277, 278]. Despite this, a comprehensive microscopic understanding of the origin, the statistics, and the effect of long-range resonances is still lacking.

As a result, recent studies have shifted focus to the properties of many-body resonances in strongly disordered regimes—regions that standard diagnostics had previously identified as many-body localized. Approaches include scaling analyses of extreme values of spectral observables [257, 258], detailed studies of large deviations in spatial correlations of the relevant degrees of freedom (such as longitudinal spin–spin correlations) [279, 280], and proxies for the probability of decorrelation from a randomly initialized state [281]. The probability distributions of these observables display heavy tails, with rare events taking values more typical of ergodic regimes. Accounting for these heavy tails shifts the estimated critical disorder strength at which the system is typically localized, predicting a new threshold $W_{\rm MBL}$ that is larger than previously thought.

In this chapter, we build upon the approach introduced in Ref. [281], previously applied to the out-of-equilibrium phase diagram of the random-field Ising model in a transverse field (also known as the Imbrie model), where the absence of diffusion at strong disorder has been rigorously established under minimal assumptions [264, 275, 276]. We extend this method to the more debated case of the random-field Heisenberg chain, a model that has been central to most numerical investigations of the MBL transition [244–247, 257, 263, 282–286]. Specifically, we study the general random-field XXZ chain and present an updated phase diagram at high energy (i.e., in the middle of the many-body spectrum) in the disorder–interaction (W, Δ) plane, accounting for the role of long-range resonances.

Building on an analogy with mean-field glassy systems, our method evaluates how often rare events occur in which the observable behaves as in the ergodic phase, caused by resonances that extend across the entire Hilbert space. The central quantity we focus on is the probability that a system initialized in a random configuration $|0\rangle$ at time t=0 is found in a configuration $|f\rangle$, located far from $|0\rangle$ in Hilbert space, at infinite time. These probabilities are estimated via the amplitudes of the propagators $|\mathcal{G}_{0f}|^2$, which are significantly easier to compute numerically. Rare resonances are identified as outliers in the probability distributions of these propagators, corresponding to pairs of resonant states separated by large distances on the Hilbert space graph.

Exploiting the analogy with classical disordered models, we introduce an auxiliary parameter β —that corresponds to the inverse temperature in the original classical model. This parameter controls the influence of extreme outliers and enables us to isolate their contribution to transport. This reveals three distinct regimes, illustrated in the phase diagram of Fig. 3.14: (i) In the ergodic phase, a system initialized in a many-body basis state $|0\rangle$ becomes delocalized over Hilbert space, exploring configurations far from the initial state; (ii) an intermediate regime where delocalization is driven by rare, disorder-

dependent long-range resonances that appear only in atypical disorder realizations; and (iii) a robust many-body localized phase, where such resonances are neither strong nor frequent enough to destabilize localization. Importantly, numerically accessible typical samples in the intermediate regime do not exhibit the system-wide resonances responsible for asymptotic delocalization. Nevertheless, our approach, guided by the analogy with mean-field glassy systems, captures the effect of these rare contributions in the large $\it L$ limit.

We investigate the spatial structure of these rare events within Hilbert space. In the MBL regime, rare resonances become increasingly short-ranged: the portion of Hilbert space accessible from the initial configuration progressively shrinks, and the spreading of the wavepacket remains confined near the initial state. In contrast, in the ergodic regime, uniform delocalization is recovered only at large distances in Hilbert space. Our results thus provide a Hilbert space-based complement to real space and spectral probes of MBL, highlighting the crucial role of rare, system-wide resonances in driving finite-size delocalization.

The emergence of heterogeneous resonant pathways facilitating delocalization is further confirmed by examining the structure of many-body eigenstates, which exhibit pronounced amplitude fluctuations in the intermediate phase, both between rare and typical disorder realizations and between distinct branches of the Hilbert space graph.

Finally, we explore these differences by visualizing the transmission pathways on the Hilbert space graph using techniques originally developed for quantum transport in mesoscopic systems [287–289]. This graphical perspective offers new insight into the destabilization of the MBL phase at finite sizes, interpreting it as the emergence of resonant transmission paths that are abundant in the ergodic regime but become increasingly rare and short-ranged deep in the localized phase.

Our results provide new numerical evidence that the introduction of interactions induces delocalization through genuinely non-perturbative mechanisms. Notably, even at very small interaction strengths Δ , the critical disorder strength separating ergodic from non-ergodic phases remains finite. This implies the existence of a broad region of the phase diagram at small but finite disorder where the addition of an infinitesimal interaction is sufficient to destroy the Anderson insulator at $\Delta=0$, and possibly even restore full ergodicity. This observation is consistent with the behavior of longitudinal correlation functions and with the recently updated phase diagram of the XXZ model at mid-spectrum energies reported in Refs. [279, 280].

This chapter is organized as follows. Sec. 3.1 reviews the non-interacting Anderson localization problem, outlining its key phenomenology and theoretical framework. Sec. 3.2 extends the discussion to interacting systems, presenting the main results that support the persistence of localization in the presence of interactions. We describe the quasilocal integrals of motion pciture that underlie the microscopic mechanism by which many-body localization (MBL) breaks ergodicity. We compare this mechanism to other forms of ergodicity breaking such as mean-field spin glasses and integrable systems. Sec. 3.3 introduces the interacting spin model central to our study: the XXZ chain in a random field, a paradigmatic model of the MBL transition. In Sec. 3.4, we review standard observables used to probe the MBL phase and discuss their limitations, along with two key mechanisms—many-body resonances and thermal avalanches—that are related to the restoration of ergodicity in the MBL phase. Building on this, Sec. 3.5 presents our main method, designed to address the shortcomings of previous observables. Sec. 3.6 reports the results obtained with this approach, characterizing finite-size regimes and the nature of many-body resonances within them. Sec. 3.7 compares rare and typical disorder realizations identified by our method, examining eigenstate properties and other observables across regimes. Finally, Sec. 3.8 summarizes our findings and suggests directions for future work. Additional technical details and supporting results are provided in four appendices.

3.1 Anderson Localization

3.1.1 The Anderson model

Experiments conducted by Feher in the mid-1950s, particularly on silicon (Si) doped with dilute phosphorus (P) impurities, demonstrated that spin excitations often remained localized rather than spreading freely [290]. In other words, instead of the spin 'flips' propagating through the material like a wave, they tended to stay near the impurity sites, as if trapped, preventing the expected movement of spin information through the lattice. This effect highlighted how disorder in the impurity energy levels hinders the transport properties of quantum systems.

To explain this, In 1958 Philip W. Anderson in his seminal work [88] studied a simple model of electrons allowed to hop between neighboring sites on a lattice, but each site possesses an energy chosen from a random distribution. In this model there is a competition between the particle's kinetic energy (hopping) and the energy mismatch caused by the disordered potential (on-site random energy). Anderson demonstrated that, if this disorder is sufficiently strong, or if the energy is close to the band edges, the particle's quantum states become spatially localized. This phenomenon is nowadays known as *Anderson localization*.

Over the decades since its appearance, the concepts and results of Anderson localization have permeated a wide range of other physical systems. It has been widely recognized as a universal wave-interference phenomenon, applying not only to electrons but also to classical waves and ultra-cold atoms [291].

The original model studied by Anderson consisted of a modified tight-binding model with on-site quenched disorder. Set on a cubic d-dimensional lattice with linear dimension L, its Hamiltonian reads

$$\hat{\mathcal{H}}_0 = -J \sum_{\langle ij \rangle} (\hat{c}_i^{\dagger} \hat{c}_j + \text{h.c.}) + \sum_i \varepsilon_i \hat{c}_i^{\dagger} \hat{c}_i , \qquad (3.1)$$

where \hat{c}_i^{\dagger} (\hat{c}_i) denotes the fermionic creation (annihilation) operator at site i, J is the hopping amplitude between the neighboring sites i and j, and ε_i represents the on-site disorder. The disorder energies ε_i are independent random variables uniformly distributed over the interval [-W/2, W/2], with zero mean and variance $W^2/12$.

The core result of Anderson is that the eigenstates of the Hamiltonian, at a given energy, can be either 'extended' or 'localized' depending on the dimensionality of the system and the parameters J and W. For the case W/J=0 the problem simplifies to free fermions on a lattice at half-filling $\hat{\mathcal{H}}_{\mathrm{FF}}$, whose eigenfunctions are plain-waves of the form

$$\phi_{\alpha}(\mathbf{r}) = \frac{1}{\sqrt{L^d}} \exp\left[\frac{2\pi i}{L} \mathbf{n}(\alpha) \cdot \mathbf{r}\right], \tag{3.2}$$

where ${\boldsymbol r}$ can take values among the L^d position vectors associated to each site, with lattice spacing a. The wavevector ${\boldsymbol n}(\alpha)=\left(n_1(\alpha),\,n_2(\alpha),\,\ldots,\,n_d(\alpha)\right)$ has components $n_k(\alpha)\in\{0,\ldots,L-1\}$. The integer $\alpha\in\{1,\ldots,L^d\}$ enumerates the eigenstates by encoding the multi-index ${\boldsymbol n}(\alpha)$. We have chosen the encoding $\alpha=1+n_1+L\,n_2+L^2\,n_3+\cdots+L^{d-1}\,n_d$ for the eigenvalue equation $\hat{\mathcal H}_{\rm FF}\,\phi_\alpha({\boldsymbol r})=E_\alpha\,\phi_\alpha({\boldsymbol r})$, with eigenenergies given by

$$E_{\alpha} = -2t \sum_{k=1}^{d} \cos\left(\frac{2\pi \, n_k(\alpha)}{L}\right). \tag{3.3}$$

In this case, the system is said to be 'delocalized' or 'extended', since a wavepacket initially prepared on site 0 at the initial time t=0 will spread under unitary evolution¹

$$|\boldsymbol{r}(t)\rangle = \hat{\mathcal{U}}(t)|\boldsymbol{r}_0\rangle = e^{-i\hat{\mathcal{H}}t}|\boldsymbol{r}_0\rangle$$
 (3.4)

A quantitative measure of this spreading is given by the mean-square displacement—where the mean is taken over lattice sites [292]

$$\overline{r^2(t)} = \sum_{j=0}^{L^d-1} |\boldsymbol{r}_j|^2 |\langle \boldsymbol{r}_j | \boldsymbol{r}(t) \rangle|^2 = \sum_{j=0}^{L^d-1} |\boldsymbol{r}_j|^2 \left| \frac{1}{\sqrt{L^d}} \sum_{\alpha} e^{-iE_{\alpha}t} \phi_{\alpha}(\boldsymbol{r}_j) \right|^2, \quad (3.5)$$

that is found to be $\overline{r^2(t)} = 2dJ^2t^2$ indicating ballistic transport in every direction. The wavepacket is found with uniform probability across all sites of the lattice after infinite time.

With sufficiently strong disorder W, the nature of the eigenstates changes: wavefunctions become exponentially localized around some fixed position \mathbf{R}_{α} in the lattice, normally denoted as the center of the α -eigenstate. In this case, the wavepacket has a probability amplitude given by

$$|\phi_{\alpha}(\mathbf{r})|^2 \sim e^{-\frac{|\mathbf{r} - \mathbf{R}_{\alpha}|}{\xi_{\text{loc}}}} \tag{3.6}$$

where $\xi_{\rm loc}$ is the disorder-dependent localization length characterizing the exponential decay. For the infinite-disorder limit, in which J/W=0, eigenstates are perfectly localized, each of them occupying a single site on the lattice, i.e. $\phi_{\alpha}(\boldsymbol{r}_i)=\delta_{i\alpha}$. Instead at zero disorder W/J=0 the extended Bloch states of Eq. (3.2) are recovered. Any intermediate case depends on the ratio J/W, and the limiting cases can be further extracted by the behavior of the localization length $\xi_{\rm loc}$:

$$\lim_{J/W \to 0} \xi_{\text{loc}} = a , \qquad (3.7)$$

and the eigenstates are localized within the position R_{α} . Whereas for

$$\lim_{J/W \to \infty} \xi_{\text{loc}} = \infty , \qquad (3.8)$$

and the eigenstates are extended, free to propagate across the entire lattice.

If we consider the hopping term acting as a perturbation then $\hat{K} = -J \sum_{\langle ij \rangle} (\hat{c}_i^{\dagger} \hat{c}_j + \text{h.c.})$. For an electron initially prepared on a site i the addition of the hopping term yields a transition probability to a new site j, that according to Fermi's Golden rule at first order of perturbation theory the transition probability is

$$\mathbb{P}_{i \to j} = 2\pi |\langle j | \hat{V} | i \rangle|^2 \delta(\varepsilon_i - \varepsilon_j) , \qquad (3.9)$$

where the matrix element $\langle j|\hat{K}|i\rangle=K_{ij}=-J$ for i and j nearest neighbors and zero otherwise. In the former case, the transition rate is significant only when the energies are nearly degenerate with respect to the energy scale set by the hopping term $|\varepsilon_i-\varepsilon_j|\lesssim J$. This condition produces a resonance between these sites and the states $|i\rangle$ and $|j\rangle$ are said to hybridize i.e. they form a quantum superposition. This can be better exemplified from the two-site Hamiltonian

$$\mathcal{H} = \begin{pmatrix} \varepsilon_i & -J \\ -J & \varepsilon_j \end{pmatrix}. \tag{3.10}$$

The eigenvalues are obtained by solving $det(\mathcal{H} - E\mathbb{I}) = 0$:

$$E_{\pm} = \frac{\varepsilon_i + \varepsilon_j}{2} \pm \sqrt{\left(\frac{\varepsilon_i - \varepsilon_j}{2}\right)^2 + J^2}.$$
 (3.11)

¹We adopt natural units throughout, setting $\hbar = 1$

The corresponding normalized eigenstates are

$$|\psi_{+}\rangle = \cos\theta |i\rangle + \sin\theta |j\rangle,$$
 (3.12)

$$|\psi_{-}\rangle = -\sin\theta \,|i\rangle + \cos\theta \,|j\rangle,\tag{3.13}$$

where the mixing angle θ is defined by

$$\tan 2\theta = \frac{2J}{\varepsilon_i - \varepsilon_j}. (3.14)$$

In the strongly resonant case $|\varepsilon_i - \varepsilon_j| \ll J$, we have $\theta \approx \pi/4$, so the eigenstates are nearly equal superpositions:

$$|\psi_{\pm}\rangle \approx \frac{1}{\sqrt{2}} \left(|i\rangle \pm |j\rangle\right),$$
 (3.15)

meaning that the wavepacket spreads equally in both sites.

In the Anderson model, the presence of disorder in the energies $\{\varepsilon_i\}$ propitiates an energy mismatch between neighboring sites. For sufficiently strong disorder $W\gg J$, resonant transitions between neighboring sites are strongly suppressed. The same holds for distant sites: tunneling processes between sites i and j separated by n lattice sites typically occur in the n-th order of perturbation theory where we expand in terms of J/W. Therefore, the transition amplitudes scale as $(J/W)^n$. On the other hand, as the quenched random energies are uncorrelated variables, counting statistics predicts that the energy difference $\delta\varepsilon = |\varepsilon_i - \varepsilon_j|$ will decay only algebraically with the distance $\delta\varepsilon \sim W/n^d$. The competition between hopping transitions and disorder strength W yields a critical value W_c that separates the ergodic phase, characterized by extended eigenstates, from the localized phase, where eigenstates are spatially confined and transport is suppressed. This critical value depends strongly on the dimensionality d.

In low spatial dimensions (d=1,2), all single-particle eigenstates become localized for any non-zero disorder. This has been rigorously proven in d=1 [293] and supported in d=2 by a renormalization group analysis [294], as well as by numerical studies [295]. Consequently, the critical disorder strength is $W_c=0$ in these cases and the localization length scales as $\xi_{\rm loc}\sim 1/W^2$. The scaling theory of localization, developed in Ref.[296] proposed a one-parameter scaling hypothesis that describes the renormalization group flow of the dimensionless conductance g(L). The associated beta function

$$\beta(g) = \frac{d \ln g(L)}{d \ln L} \,, \tag{3.16}$$

describes the scaling of the conductance with size L, and the fixed point $\beta(g)=0$ signals the metal-insulator (extended-localized) transition. This analysis found the lower critical dimension to be $D_l=2$. Consequently, for any non-zero disorder width W all eigenstates are localized for $d\leq 2$ whereas d>2 exhibit a metal-insulator transition at a non-zero value of W_c . In the latter cases, the transition can be characterized in terms of the localization length $\xi_{\rm loc}$, playing the role of the correlation length in continuous phase transitions, diverging as we approach W_c . Accordingly, coming from the localized phase the localization length $\xi_{\rm loc}$ follows the universal scaling

$$\xi_{\rm loc} \sim (W - W_c)^{-\nu}$$
 (3.17)

The best numerical estimates at d=3, for energies in the middle of the band, found the localization transition to be $W_c \simeq 16.5$ with critical exponent around $\nu=1.58(1)$ [297–305]. For higher dimensions, numerical estimates for these quantities have been performed [306, 307], they are summarized in Table 3.1

Table 3.1: Critical disorder W_c and exponent ν for dimensions d=4,5,6. Adapted from Ref. [307].

The Anderson model on the Bethe lattice [308] serves as a representative example of the infinite-dimensional limit, as discussed in previous chapters. This model is exactly solvable, and its phase diagram can be expressed as a function of the lattice connectivity $\kappa+1$; an example for $\kappa+1=3$ is shown in Fig. 3.1. In this case, the *mobility edge*—which separates localized from extended eigenstates—can be determined exactly, and the critical disorder strength W_c increases with connectivity. The corresponding critical exponent is $\nu=1/2$, in agreement with other mean-field approaches [309].

In finite-dimensional systems with d>2, a mobility edge is likewise expected, separating localized and extended states in the W-E plane. Interestingly, the critical exponent ν decreases smoothly with increasing d, showing no sign of saturating to the infinite-dimensional value $\nu=1/2$. This continuous decrease has led several authors to propose that the upper critical dimension for Anderson localization is effectively infinite [307, 310–313].

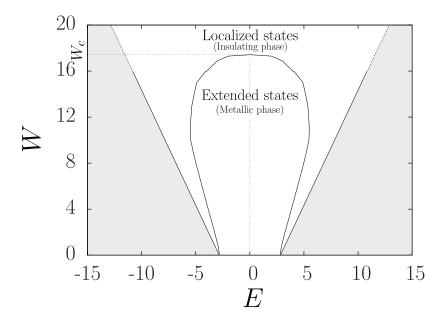


Fig. 3.1: Phase diagram for the Bethe lattice with connectivity $\kappa+1=3$. The innermost solid line indicates the mobility edge between extended and localized states, the outermost solid line being the edge of the density of states $E=\pm(2\sqrt{\kappa}+W/2)$, outside of which no eigenstates can be found (gray regions). The critical disorder width W_c at the center of the energy band E=0 is shown. Figure adapted from Ref. [308].

3.1.2 Eigenstate characterization

One useful probe to distinguish between ergodic and localized eigenstates is given by the inverse participation ratio (IPR), that provides a measure of the concentration of the wavefunction in specific sites. Its most general form is given by

$$\mathcal{I}_q^{(\alpha)} \equiv \sum_{i=0}^{L-1} |\phi_{\alpha}(\boldsymbol{r}_i)|^{2q} , \qquad (3.18)$$

In most applications, the term 'IPR' specifically refers to the case q=2. By contrast, for q=1 we have $\mathcal{I}_1^{(\alpha)}=1$, which simply reflects the normalization of the wavefunctions.

For a finite-dimensional system with linear size L, the IPR provides valuable information about the nature of the eigenstates. Focusing on a single eigenstate $|\phi_{\alpha}\rangle$, we see that:

- (i) If $|\phi_{\alpha}\rangle$ is extended, its amplitude is uniform across the lattice hence $|\phi_{\alpha}(\mathbf{r})|^2 \sim 1/L^d$, and therefore we expect $\mathcal{I}_q^{(\alpha)} = L^{-d(q-1)}$.
- (ii) If instead $|\phi_{\alpha}\rangle$ is localized, its amplitude is concentrated within the localization length $\xi_{\rm loc}$ around its center R_{α} , for these cases we expect $|\phi_{\alpha}(\boldsymbol{r})|^2 \sim O(1)$ while all the others are vanishingly small. Therefore we expect $\mathcal{I}_q^{(\alpha)} \sim O(1)$, independent of L.

Consequently, in the thermodynamic limit $L \to \infty$, the IPR is expected to vanish for extended eigenstates, whereas it remains finite for localized eigenstates.

To obtain statistically meaningful results, it is common to average $\mathcal{I}_q^{(\alpha)}$ over several eigenstates of the system within a close energy range—and over several disorder realizations. Often one studies the average IPR \mathcal{I}_q and its typical value $\mathcal{I}_q^{\mathrm{typ}}$

$$\mathcal{I}_q = \mathbb{E}\left[\mathcal{I}_q^{(\alpha)}\right] , \qquad \qquad \mathcal{I}_q = \exp \mathbb{E}\left[\ln \mathcal{I}_q^{(\alpha)}\right] , \qquad (3.19)$$

where $\mathbb{E}[\cdots]$ is taken over several disorder realizations $\{\varepsilon_i\}$ and over eigenstates whose energies lie within a narrow energy window—for example, around E=0 for the band center.

At the metal-insulator transition, \mathcal{I}_q in dimensions greater than two (d > 2), both average and typical IPRs are found to follow a continuous scaling relation

with $\tau_q = D_q(q-1)$. The cases mentioned above are recovered with $D_q = d$ for extended states, and $D_q = 0$ for localized states. The exact relationship between the two exponents τ_q and $\tau_q^{\rm typ}$ was studied in detail in Refs. [314, 315]. The quantity D_q corresponds to a fractal dimension, that has a non-trivial dependence on q and the wavefunctions are said to be multifractal [309, 313, 316], a concept we will mention briefly in the following section.

3.1.3 Spectral Properties

To gain a fundamental understanding of the localization transition, spectral theory [317, 318] provides a rigorous formal framework. For a general Hamiltonian $\hat{\mathcal{H}}$ (or in general for any self-adjoint operator) acting in a Hilbert space \mathscr{H} , the latter can be decomposed into three orthogonal components

$$\mathscr{H} = \mathscr{H}_{D} \oplus \mathscr{H}_{SC} \oplus \mathscr{H}_{AC} \tag{3.21}$$

with \mathscr{H}_D corresponding to a subspace with discrete—or pure point—spectrum $\sigma_D(\hat{\mathcal{H}})$, \mathscr{H}_{SC} corresponding to a subspace with singular continuous spectrum $\sigma_{SC}(\hat{\mathcal{H}})$, and \mathscr{H}_{AC} corresponding to a subspace with absolute continuous spectrum $\sigma_{AC}(\hat{\mathcal{H}})$. Together they form the spectrum of the operator $\hat{\mathcal{H}}$

$$\sigma(\hat{\mathcal{H}}) = \sigma_{\mathrm{D}}(\hat{\mathcal{H}}) + \sigma_{\mathrm{SC}}(\hat{\mathcal{H}}) + \sigma_{\mathrm{AC}}(\hat{\mathcal{H}})$$
(3.22)

The Anderson problem (or any other system exhibiting Anderson localization) is such that, in the thermodynamic limit, the nature of the relevant spectrum collapses into one of these three possible partitions of the Hilbert space. In other words, depending on the disorder strength, in the thermodynamic limit $L \to \infty$ the spectrum becomes:

- (i) $\sigma_D(\hat{\mathcal{H}})$: The associated eigenstates are spatially localized, and the phase is insulating—conductance is absent.
- (ii) $\sigma_{SC}(\hat{\mathcal{H}})$: The generalized eigenstates are neither fully localized nor fully extended, but 'critical'. In many physical systems these eigenstates are multifractal—a case that we won't discuss here.
- (iii) $\sigma_{AC}(\hat{\mathcal{H}})$: The generalized eigenstates are extended, and corresponding wavepackets spread over time. The phase is thus said to be metallic or ergodic.

This classification follows from the RAGE theorem, which has been rigorously proven in mathematical physics [293]. The theorem links the spectral decomposition of the Hamiltonian in the thermodynamic limit to the asymptotic infinite-time evolution of wavepackets under unitary dynamics $e^{-i\hat{\mathcal{H}}t}$.

The Resolvent Operator

This spectral information can be studied from the Green function—or resolvent operator—defined as

$$\hat{\mathcal{G}}(E) = (E\hat{\mathbb{I}} - \hat{\mathcal{H}})^{-1}$$
, (3.23)

where E is the energy at which we are evaluating the resolvent.

For finite systems of size L^d the spectrum is always discrete, and we can expand the resolvent in terms of the eigenpairs of the Hamiltonian $\hat{\mathcal{H}}$ as

$$\hat{\mathcal{G}}(E) = \sum_{\alpha} \frac{|\phi_{\alpha}\rangle \langle \phi_{\alpha}|}{E - E_{\alpha}}, \qquad (3.24)$$

with $\hat{\mathcal{H}} |\phi_{\alpha}\rangle = E_{\alpha} |\phi_{\alpha}\rangle$. In the position representation this becomes

$$\mathcal{G}(\boldsymbol{r}, \boldsymbol{r}'; E) = \langle \boldsymbol{r} | \hat{\mathcal{G}} | \boldsymbol{r}' \rangle = \sum_{\alpha} \frac{\phi_{\alpha}(\boldsymbol{r}) \phi_{\alpha}^{*}(\boldsymbol{r}')}{E - E_{\alpha}}.$$
(3.25)

The function $\mathcal{G}(\boldsymbol{r},\boldsymbol{r}';E)$ is analytic in E, except for simple poles located at $E=E_{\alpha}$. In the thermodynamic limit, the spectrum becomes continuous in the ergodic phase, and Eq. (3.25) changes accordingly. For instance, in the disorder-free tight-binding model for $L\to\infty$, the Green function is expressed in terms of the continuous momentum \boldsymbol{k}

$$\mathcal{G}(\boldsymbol{r}, \boldsymbol{r}'; E) = \int \frac{d^d k}{(2\pi)^d} \frac{\phi_{\boldsymbol{k}}(\boldsymbol{r})\phi_{\boldsymbol{k}}^*(\boldsymbol{r}')}{E - E(\boldsymbol{k})}.$$
 (3.26)

More generally, when the spectrum is continuous, the discrete poles of \mathcal{G} merge into a continuum with support on the real axis. In the example above, this continuum is determined by the values of the dispersion relation $E(\mathbf{k})$ appearing in the denominator of Eq. (3.26). Conversely, in the localized phase, the spectrum remains discrete in the thermodynamic limit, and the singular behavior of \mathcal{G} persists only through the isolated poles at energies E_{α} .

As a result, the two phases can be distinguished by examining the singular behavior of the Green function through a limiting procedure. This involves analytically continuing

G(r, r'; E) into the complex plane by adding a small imaginary part to the energy, controlled by a parameter η .

For an absolute continuous spectrum, taking the limit $\eta \to 0$ causes the analytically continued Green function to develop a branch cut along the real axis—or, in the case of a singular continuous spectrum, a natural boundary. The two analytic continuations across this branch cut correspond to

$$\mathcal{G}^{r}(\boldsymbol{r}, \boldsymbol{r}'; E) = \lim_{\eta \to 0^{+}} \mathcal{G}(\boldsymbol{r}, \boldsymbol{r}'; E - i\eta) ,$$

$$\mathcal{G}^{a}(\boldsymbol{r}, \boldsymbol{r}'; E) = \lim_{\eta \to 0^{+}} \mathcal{G}(\boldsymbol{r}, \boldsymbol{r}'; E + i\eta) ,$$
(3.27)

where \mathcal{G}^r and \mathcal{G}^a are the 'retarded' and 'advanced' Green functions, respectively. Their respective operators are related as $(\hat{\mathcal{G}}^r)^{\dagger} = \hat{\mathcal{G}}^a$.

Both quantities in Eq. (3.27) are equal if no singular point is encountered in the limit $\eta \to 0^+$. However, for the continuous spectrum, taking the difference between retarded and advanced Green functions quantifies the discontinuous jump across the branch cut. While in the discrete spectrum this discontinuity is not present. Instead, the Green function simply diverges at each pole, and the limit is ill-defined precisely at those points. Away from the poles, however, the difference is zero, as no singularity is encountered with $\eta \to 0$. This difference can be expressed simply as the limit $\eta \to 0^+$ of the imaginary part of the analytically continued Green function:

$$\Delta \mathcal{G}(\mathbf{r}, \mathbf{r}'; E, \eta) \equiv \mathcal{G}^{r}(\mathbf{r}, \mathbf{r}'; E) - \mathcal{G}^{a}(\mathbf{r}, \mathbf{r}'; E)$$

$$= 2i \operatorname{Im} \mathcal{G}(\mathbf{r}, \mathbf{r}'; E - i\eta) .$$
(3.28)

This imaginary part, in turn, is directly related to the *local density of states* (LDOS) $\rho(\mathbf{r}; E)$. To see this, consider Eq. (3.25) in the thermodynamic limit $L \to \infty$. By analytically continuing the Green function as $E \to E - i\eta$, setting $\mathbf{r}' = \mathbf{r}$, and then taking the limit $\eta \to 0^+$, the connection to the LDOS becomes evident.

This follows from the Sokhotski–Plemelj theorem (also known as Dirac's identity), which states:

$$\lim_{\eta \to 0^+} \left(\frac{1}{x \pm i\eta} \right) = \mathscr{P}\left(\frac{1}{x} \right) \mp i\pi \delta(x), \tag{3.29}$$

with \mathcal{P} being the Cauchy principal value.

For instance, in the localized phase—where Eq. (3.25) remains valid even in the thermodynamic limit—the local density of states is found from

$$\rho(\mathbf{r}; E) = \sum_{\alpha} |\phi_{\alpha}(\mathbf{r})|^2 \delta(E - E_{\alpha}) = \lim_{\eta \to 0^+} \left\{ \frac{1}{\pi} \operatorname{Im} \mathcal{G}(\mathbf{r}, \mathbf{r}; E - i\eta) \right\} . \tag{3.30}$$

Consequently, the local density of states contains all the relevant information of the localization transition, and can serve as an order parameter for the transition.

The local density of states (LDOS) quantifies the spectral weight at a given energy around a specific position r. In the localized phase, this spectral weight is concentrated at discrete energies corresponding to isolated poles of the Green function, and the LDOS vanishes almost everywhere except at those energies. Only eigenstates that have significant amplitude at position r contribute to the LDOS at that site. These are the eigenstates whose localization centers \mathbf{R}_{α} lie within a localization length of r—that is, those for which $|\phi_{\alpha}(r)|^2$ is non-negligible.

In contrast, in the ergodic phase, eigenstates are delocalized and spread uniformly over the system, with amplitudes typically of order 1/N at each site. The spectrum becomes

continuous and has support over an interval on the real axis—possibly the entire axis. As a result, the LDOS is finite for any position r within this spectral interval, since a finite measure of delocalized states contributes at every site.

Moreover, the local density of states is a random variable because it depends on the specific realization of the quenched disorder $\{\varepsilon_i\}$. Consequently, the order parameter of the transition is formally described by the probability distribution function $\mathcal{P}(\rho(r; E))$, defined over both lattice sites and disorder realizations $\{\varepsilon_i\}$.

In the localized phase, the probability distribution function (PDF) of the LDOS is not a smooth function. The LDOS consists of delta peaks located at discrete (random) energies E_{α} . As a result, for a fixed energy E, the quantity $\rho({\bf r};E)$ is typically zero—since E almost surely does not coincide with any of the eigenenergies—and diverges at isolated points where $E=E_{\alpha}$. The PDF is therefore ill-defined in the usual sense and becomes singular in the localized phase. This pathological behavior mirrors the behavior of the Green function in the limit $\lim_{\eta\to 0} \operatorname{Im} \mathcal{G}({\bf r},{\bf r},E-i\eta)$. By contrast, in the ergodic phase, the PDF of the local density of states is well-defined and smooth.

An equivalent characterization of the transition can be given in terms of the *typical* value of the local density of states,

$$\rho_{\text{tvp}}(E) = \exp\left(\mathbb{E}\left[\ln \, \rho(\boldsymbol{r}; E)\right]\right) \tag{3.31}$$

where the average $\mathbb{E}[\cdots]$ is taken over lattice sites and disorder realizations. This quantity is well-defined for both phases, yielding zero at the localized phase $(W > W_c)$ and nonzero in the extended phase $(W < W_c)$, providing yet another order parameter to the localization transition:

$$\rho_{\text{typ}}(\boldsymbol{r};E) \neq 0 , \qquad \text{for} \quad W < W_c$$

$$\rho_{\text{typ}}(\boldsymbol{r};E) = 0 , \qquad \text{for} \quad W > W_c .$$

$$(3.32)$$

This quantity should be contrasted with the average global density of states, defined as

$$\rho(E) = \mathbb{E}\left[\frac{1}{L^d} \sum_{j=0}^{L^d-1} \rho(\boldsymbol{r}_j; E)\right]$$
(3.33)

where $\mathbb{E}[\cdots]$ is the average over disorder realizations.

In the ergodic phase, the two quantities coincide, $\rho(E) = \rho_{\rm typ}(E)$. In the localized phase, however, the typical value $\rho_{\rm typ}(E)$ vanishes, while the distribution $\mathcal{P}(\rho(\boldsymbol{r};E))$ develops fat tails arising from rare, large values at the poles $E=E_{\alpha}$. As a result, the global density of states $\rho(E)$ remains finite even in the localized phase, rendering it unsuitable for identifying the transition.

The Locator Expansion—and the Self-Energy

In fact, the Green function—and its spectral properties—played a central role in Anderson's original work, where it formed the basis of the celebrated locator expansion used to probe localization [88]. We can compute $\mathcal{G}(\boldsymbol{r},\boldsymbol{r};E)$ by separating the Hamiltonian into kinetic, $\hat{K}=-J\sum_{\langle ij\rangle}(\hat{c}_i^{\dagger}\hat{c}_j+\text{h.c})$, and potential terms $\hat{V}=\sum_i\varepsilon_i\hat{c}_i^{\dagger}\hat{c}_i$. Equivalently, we can separate the calculation of the resolvent in a bare term $\hat{\mathcal{G}}_V$ —containing just the random potential—and the full resolvent $\hat{\mathcal{G}}$. We omit the energy dependency to simplify the notation and write:

$$\hat{\mathcal{G}} = \frac{1}{E\mathbb{I} - \hat{\mathcal{H}}}, \qquad \qquad \hat{\mathcal{G}}_V = \frac{1}{E\mathbb{I} - \hat{V}}. \tag{3.34}$$

From these expressions, it is straightforward to obtain the operator identity (Dyson's equation):

$$\hat{\mathcal{G}} = \hat{\mathcal{G}}_V + \hat{\mathcal{G}}_V \hat{K} \hat{\mathcal{G}} . \tag{3.35}$$

Taking the expectation value over the site at position r

$$\mathcal{G}(\mathbf{r}, \mathbf{r}) = \langle \mathbf{r} | \hat{\mathcal{G}} | \mathbf{r} \rangle = \mathcal{G}_V(\mathbf{r}, \mathbf{r}) + \langle \mathbf{r} | \hat{\mathcal{G}}_V \hat{K} \hat{\mathcal{G}} | \mathbf{r} \rangle . \tag{3.36}$$

We can keep expanding the second term on the right hand side, to obtain the expansion

$$\mathcal{G}(\boldsymbol{r}, \boldsymbol{r}) = \mathcal{G}_{V}(\boldsymbol{r}, \boldsymbol{r}) + \sum_{i} \mathcal{G}_{V}(\boldsymbol{r}, \boldsymbol{r}) K(\boldsymbol{r}, \boldsymbol{r}_{i}) \mathcal{G}_{V}(\boldsymbol{r}_{i}, \boldsymbol{r}_{i})
+ \sum_{i,j} \mathcal{G}_{V}(\boldsymbol{r}, \boldsymbol{r}) K(\boldsymbol{r}, \boldsymbol{r}_{i}) \mathcal{G}_{V}(\boldsymbol{r}_{i}, \boldsymbol{r}_{i}) K(\boldsymbol{r}_{i}, \boldsymbol{r}_{j}) \mathcal{G}_{V}(\boldsymbol{r}_{j}, \boldsymbol{r}_{j}) + \dots$$
(3.37)

This is just a perturbation expansion over the hopping parameter—known as the locator expansion. The zeroth order corresponds to no hopping $\mathcal{G}^{(0)}(\boldsymbol{r},\boldsymbol{r})=1/(E-\varepsilon_0)$, the second order term corresponds to a round trip of two steps $0 \to i \to 0$, given by

$$\mathcal{G}^{(2)}(\boldsymbol{r},\boldsymbol{r}) = \sum_{i} \frac{J_{0i}^{2}}{(E - \varepsilon_{0})(E - \varepsilon_{i})(E - \varepsilon_{0})},$$
(3.38)

with J_{0i} the hopping term with a reintroduced site dependence. We repeat this order iteratively up to the desire order of perturbation theory. We can group together all the 'irreducible paths' i.e. those paths that cannot be split into two at site 0—for example $0 \to i \to j \to i \to 0$ at 4th order. The group of all irreducible contributions are collected in the *self-energy* $\Sigma(E, \mathbf{r})$ associated to the position \mathbf{r} . As a result, Dyson's equation resums all reducible diagrams, and we recover the recursive form

$$\mathcal{G}(\boldsymbol{r}, \boldsymbol{r}; E) = \mathcal{G}_{V}(\boldsymbol{r}, \boldsymbol{r}; E) + \mathcal{G}_{V}(\boldsymbol{r}, \boldsymbol{r}; E) \Sigma(\boldsymbol{r}, E) \mathcal{G}(\boldsymbol{r}, \boldsymbol{r}; E) . \tag{3.39}$$

or equivalently

$$\mathcal{G}(\boldsymbol{r}, \boldsymbol{r}; E) = \left(\mathcal{G}_V(\boldsymbol{r}, \boldsymbol{r}; E)^{-1} - \Sigma(E, \boldsymbol{r})\right)^{-1}$$
(3.40)

where we have reinserted explicitly all relevant dependencies. Using this last equation and the explicit form of the resolvent in Eq. (3.23) we get

$$\mathcal{G}(\boldsymbol{r}, \boldsymbol{r}; E - i\eta) = (E - \varepsilon_0 - i\eta - \Sigma(\boldsymbol{r}, E - i\eta))^{-1}$$
(3.41)

where ε_0 is the on-site energy associated to the site at position r.

Computing the self-energy $\Sigma(r, E)$ is particularly challenging, as it encodes the complex influence of the disordered environment on a given site. Because of this various approximations —such as truncating the locator expansion or using self-consistent methods—are employed to estimate its effects in a tractable way [319].

Its imaginary part is related to the inverse characteristic time for a particle to escape from the evaluated site into the rest of the system. In localized phases, this escape rate vanishes, while in ergodic phases it remains finite, this will be made explicit seen in the next section.

3.1.4 Relationship with the time evolution

Anderson localization is a dynamical phase of matter, as such the properties aforementioned can be understood in terms of its dynamics. In fact, the problem initially considered by

Anderson [88] was to compute the return probability $\mathcal{R}(t)$ for a particle initially prepared in site $|\mathbf{r}_0\rangle$ at t=0. This quantity is given by

$$\mathscr{R}(t) = |\langle \boldsymbol{r}_0 | \boldsymbol{r}(t) \rangle|^2 , \qquad (3.42)$$

with

$$|\boldsymbol{r}(t)\rangle = \hat{\mathcal{U}}(t) |\boldsymbol{r}_0\rangle = e^{-i\hat{\mathcal{H}}t} |\boldsymbol{r}_0\rangle = \sum_{\alpha} e^{-iE_{\alpha}t} \langle \phi_{\alpha} | \boldsymbol{r}_0 \rangle |\phi_{\alpha}\rangle .$$
 (3.43)

Using Cauchy's integral formula to express the factor $e^{-iE_{\alpha}t}$ as

$$e^{-iE_{\alpha}t} = \frac{1}{2\pi i} \oint_{\mathscr{C}} dz \frac{e^{-izt}}{z - E_{\alpha}} , \qquad (3.44)$$

with $\mathscr C$ is a contour encircling E_α counterclockwise. Then

$$\langle \boldsymbol{r}_{0} | \boldsymbol{r}(t) \rangle = \frac{1}{2\pi i} \oint_{\mathscr{C}} dz \, e^{-izt} \sum_{\alpha} \frac{|\phi_{\alpha}(\boldsymbol{r}_{0})|^{2}}{z - E_{\alpha}}$$

$$= \frac{1}{2\pi i} \oint_{\mathscr{C}} dz \, e^{-izt} \mathcal{G}(\boldsymbol{r}_{0}, \boldsymbol{r}_{0}; z) ,$$
(3.45)

where the sum over eigenstates has been identified with the diagonal element of the resolvent (Green's function) given in Eq. (3.25) with $r' = r = r_0$. In other words, the resolvent can be interpreted as the complex Laplace transform of the time evolution operator.

Setting $z=E-i\eta$, we evaluate the contour integral by enclosing the real axis using a counterclockwise rectangular contour. For the negative oriented segment (upper half-plane) and the positive oriented one (lower half-plane), we take the analytic continuations $E+i\eta$ and $E-i\eta$, respectively.

$$\langle \boldsymbol{r}_{0} | \boldsymbol{r}(t) \rangle = \frac{1}{2\pi i} \int_{-\infty}^{\infty} dE \, e^{-iEt} \left[e^{-\eta t} \mathcal{G}(\boldsymbol{r}_{0}, \boldsymbol{r}_{0}; E - i\eta) - e^{\eta t} \mathcal{G}(\boldsymbol{r}_{0}, \boldsymbol{r}_{0}; E + i\eta) \right]$$

$$(\text{taking } \eta \to 0^{+})$$

$$= \frac{1}{\pi} \int_{-\infty}^{\infty} dE \, e^{-iEt} \operatorname{Im} \mathcal{G}^{r}(\boldsymbol{r}_{0}, \boldsymbol{r}_{0}; E)$$

$$(3.46)$$

Using Eq. (3.41), we can split the self-energy into imaginary and real parts $\Sigma(\boldsymbol{r}_0, E - i\eta) = \text{Re}\Sigma(\boldsymbol{r}_0, E, \eta) + i\text{Im}\Sigma(\boldsymbol{r}_0, E, \eta)$, with $\text{Im}\Sigma(\boldsymbol{r}_0, E, \eta) > 0$. We obtain the retarded Green function of Eq. (3.45) by simply taking the limit $\eta \to 0^+$,

$$\mathcal{G}^{r}(\boldsymbol{r}_{0},\boldsymbol{r}_{0};E) = (E - \varepsilon_{0} - \operatorname{Re}\Sigma(\boldsymbol{r}_{0},E) - i\operatorname{Im}\Sigma(\boldsymbol{r}_{0},E))^{-1}.$$
(3.47)

Finally, its imaginary part can be obtained as

$$\operatorname{Im} \mathcal{G}^{r}(\boldsymbol{r}_{0}, \boldsymbol{r}_{0}; E) = \frac{\operatorname{Im} \Sigma(\boldsymbol{r}_{0}, E)}{[E - \varepsilon_{0} - \operatorname{Re} \Sigma(\boldsymbol{r}_{0}, \boldsymbol{r}_{0}; E)]^{2} + [\operatorname{Im} \Sigma(\boldsymbol{r}_{0}, E)]^{2}},$$
(3.48)

and inserting in Eq. (3.46) we recognize the expression as the inverse Fourier transform of a Cauchy distribution, yielding

$$\langle \boldsymbol{r}_0 | \boldsymbol{r}(t) \rangle \simeq \exp \left(i(\varepsilon_0 + \Sigma(\boldsymbol{r}_0, E)) \right) \exp \left(-\operatorname{Im}\Sigma(\boldsymbol{r}_0, E) t \right)$$
 (3.49)

$$\mathscr{R}(t) = |\langle \boldsymbol{r}_0 | \boldsymbol{r}(t) \rangle|^2 \propto \exp\left(-2\operatorname{Im}\Sigma(\boldsymbol{r}_0, E) t\right)$$
(3.50)

where it is clear that as $t \to \infty$ the return probability goes to zero, as expected for extended states. We have assumed that the self-energy is weakly dependent on E—that is not

problematic in the metallic phase. Here it is clear that the imaginary part of the self-energy is related to the characteristic time of decay of the return probability, and thus related to the escape rate.

We consider the strong localized regime $W \gg J$, where we can truncate the locator expansion at first order, as a result:

$$\mathcal{G}(\mathbf{r}_0, \mathbf{r}_0; E - i\eta) = (E - \varepsilon_0 - i\eta)^{-1}. \tag{3.51}$$

After taking the imaginary part we obtain:

$$\operatorname{Im} \mathcal{G}(\boldsymbol{r}_0, \boldsymbol{r}_0; E - i\eta) = \frac{\eta}{(E - \varepsilon_0)^2 + \eta^2}, \qquad (3.52)$$

taking the limit $\eta \to 0^+$ (along with the prefactor $1/\pi$ of Eq. (3.46)) we recognize the Lorentzian representation of the Dirac delta function, thus

$$\langle \mathbf{r}_0 | \mathbf{r}(t) \rangle = \int_{-\infty}^{\infty} dE \, e^{-iEt} \delta(E - \varepsilon_0) = \exp(-i\varepsilon_0 t)$$
 (3.53)

$$\mathscr{R}(t) = |\langle \mathbf{r}_0 | \mathbf{r}(t) \rangle|^2 = \exp(-2i\varepsilon_0 t)$$
(3.54)

where $t\to\infty$ gives $\mathscr{R}(t)=1$, consistent with completely localized states in the insulating phase.

In the limit $\eta \to 0^+$, the imaginary part of the resolvent is directly related to the imaginary part of the self-energy, which determines the characteristic decay time of excitations in the system. This connection highlights the deep relationship between spectral properties and real-time dynamics. Since the resolvent is the Laplace transform of the time evolution operator, it encodes energy-resolved information about dynamical processes, such as transitions between different states of the system.

3.2 Many-body localization (MBL)

3.2.1 From Anderson to Many-Body Localization

The problem of localization becomes more intricate when interactions are introduced. Mott demonstrated that coupling electrons to a phonon bath can induce a finite conductivity through phonon-assisted hopping [320]. However, in this scenario, the phonon bath remains unaffected by the localized electrons. This raised the question: Could electron-electron interactions play a role analogous to phonon-electron interactions and thermalize an otherwise localized phase? This question was already explored in the 1980s by Fleishman and Anderson [321], who, using perturbative analysis and general arguments, suggested that at sufficiently low temperatures and strong disorder localization could still persist in systems with short-range interactions.

Intuitively, the presence of interactions might facilitate resonant events and restore a finite conductance. Consider the Anderson model of Eq. (3.1) with the addition of interactions:

$$\hat{\mathcal{H}} = -J \sum_{\langle ij \rangle} (\hat{c}_i^{\dagger} \hat{c}_j + \text{h.c.}) + \sum_i \varepsilon_i \hat{n}_i + \Delta \sum_{\langle ij \rangle} \hat{n}_i \hat{n}_j , \qquad (3.55)$$

where \hat{n}_i is the number operator $\hat{n}_i = \hat{c}_i^{\dagger} \hat{c}_i$.

In the Anderson problem, the energy mismatch relevant for an electron hopping between sites i and j is given by $\delta \varepsilon = |\varepsilon_i - \varepsilon_j|$. When this mismatch is much larger than the hopping constant $\delta \varepsilon \gg J$ the transition is strongly suppressed.

In the presence of interactions, however, the energy of an electron at each site is modified by the occupations of its neighboring sites. The effective energy at site i becomes

$$\varepsilon_i^{\text{eff}} = \varepsilon_i + \Delta \sum_{j \in \partial i} n_j ,$$
 (3.56)

where n_j is the occupation number at site j (with $\hat{n}_j |j\rangle = n_j |j\rangle$).

The new energy shift introduced by Δ can partially compensate the bare energy mismatch $\delta \varepsilon$ between the two sites i and j. If site i has occupied neighbors, its energy is raised by Δ per occupied neighbor, while the energy of j is shifted according by its own local neighborhood. As a result, the effective energy mismatch $|\varepsilon_i^{\rm eff} - \varepsilon_j^{\rm eff}|$ may become smaller (or larger), depending on the occupations of the surrounding sites. A hopping process that was off-resonant in the single-particle case can thus become resonant once interaction-induced energy shifts are included, enabling the formation of many-body resonances.

The stability of the Anderson insulator under the effect of small interactions was first thoroughly investigated by the breakthrough of Basko, Aleiner, and Altshuler [89] (see also Ref. [90]). The interacting problem was formulated in Fock space: rather than working in real space, they represented each many-body configuration as a site on a high-dimensional graph, where edges correspond to interaction-induced transitions (or 'channels') between configurations. In this picture, the interacting problem becomes analogous to Anderson localization on this Fock-space graph. They found that the connectivity of this graph—i.e., the number of available paths for transitions—grows with the system's energy (or 'temperature'). Using this analogy, they identified a critical energy above which the system delocalizes, while below it the system remains localized. This provided the first theoretical evidence for the existence of a many-body localized (MBL) phase, in which localization survives despite interactions.

3.2.2 Perturbative Perspective on MBL

Basko et al. [89] used Keldysh perturbation theory and focused on a class of diagrams known as rainbow diagrams (no loops), which can be resummed within the Self-Consistent Born Approximation (SCBA). Neglecting the real part of the self-energies leads to a simplified version called Imaginary SCBA (ImSCBA). This approximation is conceptually related to the Forward Scattering Approximation used in single-particle Anderson problems [322, 323]. These approximations are known to provide an upper bound on the critical disorder strength for delocalization: if SCBA predicts localization, the system is certainly localized, but SCBA may predict delocalization even when the system is still localized, serving as good probes for the single-particle setting.

However, because SCBA ignores loop diagrams and treats only the most resonant decay paths, its predictions may not capture non-perturbative effects that could destabilize many-body localization. In particular, rare 'Griffiths regions'—regions within the system for which the disorder strength is anomalously weak—or higher-order processes could eventually induce delocalization even when perturbative calculations predict stability [324, 325]. For this reason, while Basko et al.'s work provided the first analytic evidence for the many-body localized phase, modern perspectives view their result as suggestive rather than definitive: the true stability and boundaries of MBL likely require going beyond perturbation theory. These non-perturbative effects are the main interest of the present chapter.

Most studies of many-body localization (MBL) focus on one-dimensional systems (d=1) because localization is both conceptually simpler and computationally more tractable there. In d=1, even infinitesimal disorder is sufficient to localize non-interacting particles [88], providing a natural starting point for understanding MBL. Moreover, numerical techniques

such as exact diagonalization and matrix product state methods can capture the essential physics in 1d, but become severely limited in higher dimensions due to the exponential growth of the Hilbert space [91–94].

However, even in d=1 the ultimate stability of the MBL phase remains unsettled. Several works have suggested that the rare, weakly disordered Griffiths regions could eventually destabilize localization at very long times or in the thermodynamic limit via an avalanche instability [100, 243, 248–254], that we will explain in the following sections. This leaves open important questions about whether true MBL can survive indefinitely even in d=1. Because of this unresolved issue, the situation in higher dimensions (d>1) is even less clear: while some authors have argued that MBL might persist in 2d [326, 327], others expect that the same destabilizing mechanisms are even stronger at higher dimensions. For this reason, most theoretical and numerical analyses of the MBL transition are carried out in d=1, where the phenomenology is best understood and the open questions can be most sharply addressed.

3.2.3 Contrasting MBL with Other Forms of Ergodicty Breaking

Unlike standard phase transitions, MBL is not associated with any spontaneous symmetry breaking. It shares several phenomenological features with glassy systems—both classical and quantum—such as ergodicity breaking and slow dynamics. However, the microscopic mechanisms responsible for these effects in MBL are fundamentally different from those underlying conventional glassiness [328].

Spin glasses arise in disordered systems where quenched randomness and competing (frustrated) interactions prevent the system from selecting a unique global energy minimum [35, 329]. This leads to ergodicity breaking at low temperatures: the system becomes trapped in one of many possible configurations and fails to fully explore phase space on accessible timescales.

In mean-field models, this behavior is often described in terms of a rugged free-energy landscape with an extensive number of metastable states separated by barriers that grow with system size. MBL, by contrast, is an intrinsically quantum phenomenon. It occurs in highly excited states, where the discreteness of the many-body spectrum and quantum interference prevent resonant energy exchange between configurations [91, 94].

Glassy dynamics are typically enabled and stabilized by coupling to a cold thermal bath: the bath provides the fluctuations needed for the system to gradually relax across energy barriers, and this slow relaxation fundamentally depends on the presence of such an environment. The stability of glassy phases generally increases with dimensionality; for instance, mean-field spin glasses in high dimensions exhibit replica symmetry breaking, reflecting a highly fragmented configuration space [29–35].

In contrast, many-body localized (MBL) phases are intrinsically fragile. They generally appear only in low dimensions and require strict isolation from any external bath [248, 253–261]. Even weak coupling to the environment can induce delocalization, restoring transport and ergodicity by allowing energy exchange and dephasing between localized degrees of freedom.

These contrasting responses to dimensionality and environmental coupling provide a sharp criterion to distinguish glasses from MBL systems: glasses rely on frustration and survive in contact with a bath; MBL relies on quantum interference and collapses upon coupling to a bath [328]. Despite these differences, hybrid models, like the Quantum Random Energy Model (QREM) display features of both. In the dynamical phase diagram of the QREM an MBL phase is known to emerge [330, 331], while a glassy phase is realized in equilibrium at low temperatures [332, 333].

Other mechanisms for ergodicity breaking also arise in quantum systems — for example, in integrable models [334–336], systems hosting many-body scars [337, 338], or in those exhibiting Hilbert-space fragmentation [339, 340]. However, these types of non-ergodic behavior are typically fragile: small generic perturbations break integrability, destabilize scars, or reconnect fragmented sectors, ultimately restoring thermalization.

By contrast, many-body localization represents a more robust form of ergodicity breaking. Here, quenched disorder dynamically generates an extensive set of quasi-local conserved quantities—an emergent integrability [91–94, 239–242]—which protects non-ergodic behavior even far from equilibrium. This robustness is what allows MBL to persist under conditions where other non-ergodic phenomena would typically thermalize.

3.2.4 Integrable systems

Integrability is another of the main sources of ergodicity breaking. A classical system with N degrees of freedom s_i is said to be Liouville integrable if there exist N functionally independent conserved quantities I_k that are in involution—that is, their mutual Poisson brackets vanish:

$$\{I_i, I_j\}_{PB} = 0 \qquad \forall i, j \in 1, \dots, N .$$
 (3.57)

We have already encountered the most trivial example in the introductory chapter of this thesis: the Hamiltonian is always in involution with itself, making it a constant of motion (i.e., the total energy). Building on this, we formulated the ergodic hypothesis, which in its microcanonical form states that the long-time dynamics of a system initially prepared with energy E_0 will explore all microstates compatible with E_0 in phase-space.

Put simply, the long-time average of an observable $\mathcal{O}(t)$ should coincide with its ensemble average. In this case, using the microcanonical measure, this reduces to the statement:

$$\overline{\mathcal{O}} = \lim_{t \to \infty} \frac{1}{t} \int_{\tau_R}^{\tau_R + t} dt' \, \mathcal{O}(\{s_i(t')\}) = \langle \mathcal{O} \rangle = \int \prod_{i=1}^N ds_i \, \varrho(\{s_i\}, E_0) \, \mathcal{O}(\{s_i\}) \,, \tag{3.58}$$

with the measure of the microcanonical ensemble given by

$$\varrho(\{s_i\}, E_0) = \frac{\delta(\mathcal{H}(\{s_i\}) - E_0)}{\int \prod_{i=1}^N ds_i \ \delta(\mathcal{H}(\{s_i\}) - E_0)} \ . \tag{3.59}$$

The presence of more integrals of motion impose extra constraints on the dynamics: the system's trajectory is confined to the intersection of all constant- I_k hypersurfaces. Consequently, the system cannot sample all microstates with energy E_0 ; instead, it explores only a restricted manifold determined by all conserved quantities. In contrast, a non-integrable (chaotic) system typically has fewer conserved quantities; its trajectories densely explore allowed phase-space regions, leading to ergodic behavior over long times.

One might be tempted to extend this notion directly to quantum systems via Dirac's canonical quantization, by mapping the classical conserved quantities I_k to operators \hat{I}_k and replacing Poisson brackets with commutators, proposing a definition where $[\hat{I}_i, \hat{I}_j] = 0$. However, this straightforward translation, is flawed as it fails to provide any meaningful information or to separate quantum models into distinct classes, which is a crucial requirement for a useful definition of integrability [341, 342].

The fundamental issue is that any quantum model associated with a finite-dimensional Hilbert space \mathscr{H} (e.g., systems involving spins) trivially possesses a maximal set of mutually commuting operators. This is guaranteed by the spectral theorem, which states that any hermitian Hamiltonian (\mathcal{H}) is diagonalizable. From its $\dim(\mathscr{H})$ orthogonal state vectors

(eigenstates) $|\psi_{\alpha}\rangle$, one can construct projection operators $\hat{I}_{\alpha} = |\psi_{\alpha}\rangle \langle \psi_{\alpha}|$. This set of projectors inherently constitutes a maximal independent commuting set, as they are diagonal in the Hamiltonian's eigenbasis, and their eigenvalues are sufficient to uniquely specify a state in Hilbert space. Therefore, if this naive translation of integrability were to be accepted, virtually every quantum system would be deemed 'integrable' rendering the definition meaningless for classification. This highlights the ongoing debate and the need for more nuanced definitions of quantum integrability [342].

The standard conception of quantum integrability arises in the study of models solvable via the Yang–Baxter equation—regarded as a sufficient yet not necessary condition for integrability [342]. These systems possess an extensive number of conserved quantities that can be generated systematically from the so-called transfer matrix, ensuring exact solvability through the Bethe ansatz [343]. Crucially, the conserved operators in Yang–Baxter integrable models are typically non-local [344], involving combinations of degrees of freedom across the entire system, and are associated with algebraic structures like quantum groups and R-matrices [345] rather than simple projection operators.

The many-body localized (MBL) phase is often regarded as an emergent integrable phase, in the sense that it is characterized by an extensive set of quasi-local integrals of motion (LIOMs or '\$\ell\$-bits'\$) that constrain the dynamics and prevent thermalization [91–94, 239–242]. The main difference with respect to standard quantum integrable systems is that LIOMs are exponentially localized in real space and closely tied to the microscopic degrees of freedom, in stark contrast to the non-local conserved quantities generated by the Yang–Baxter algebra. Whereas Yang–Baxter integrable models owe their solvability to fine-tuned algebraic structures—R-matrices, transfer matrices, and Bethe ansatz methods—that are typically destroyed by even infinitesimal perturbations [346, 347], the quasi-local conservation laws in MBL remain robust: local perturbations merely renormalize the LIOMs without destabilizing the non-ergodic phase [242, 275]. This robustness highlights a key distinction: while Yang–Baxter systems represent isolated solvable points in parameter space, MBL constitutes a stable dynamical phase of matter whose non-ergodic behavior survives generic deformations of the Hamiltonian.

3.2.5 Quasilocal Integrals of Motion

In order to understand the emergent integrability of the MBL phase, it is useful to first revisit the non-interacting Anderson problem. In one dimension with periodic boundary conditions, the Hamiltonian can be written as

$$\hat{\mathcal{H}} = -J \sum_{i=1}^{L} \left(\hat{c}_i^{\dagger} \hat{c}_{i+1} + \text{h.c.} \right) + \sum_{i=1}^{L} \varepsilon_i \, \hat{n}_i, \tag{3.60}$$

with $\hat{n}_i = \hat{c}_i^\dagger \hat{c}_i$ the local number operator in the site (real-space) basis.

The corresponding single-particle eigenproblem is

$$\hat{\mathcal{H}} |\phi_{\alpha}\rangle = \epsilon_{\alpha} |\phi_{\alpha}\rangle , \quad \text{with} \quad |\phi_{\alpha}\rangle = \sum_{i=1}^{L} \phi_{\alpha}(i) |i\rangle ,$$
 (3.61)

where $\alpha = 1, ..., L$ labels the single-particle eigenstates (Anderson orbitals) and $\phi_{\alpha}(i)$ are normalized real-space amplitudes.

Since $\hat{\mathcal{H}}$ is quadratic in the fermion operators, diagonalizing its single-particle form also diagonalizes the full many-body (non-interacting) problem. This is achieved by a *unitary* change of basis from the site basis $\{|i\rangle\}$ to the Anderson-orbital basis $\{|\phi_{\alpha}\rangle\}$. As a result, the

Hamiltonian is readily diagonalizable by introducing new fermionic creation/annihilation operators

$$\hat{b}_{\alpha} = \sum_{i=1}^{L} \phi_{\alpha}(i) \, \hat{c}_{i}, \qquad \hat{b}_{\alpha}^{\dagger} = \sum_{i=1}^{L} \phi_{\alpha}^{*}(i) \, \hat{c}_{i}^{\dagger},$$
 (3.62)

with the inverse relations

$$\hat{c}_i = \sum_{\alpha=1}^L \phi_{\alpha}^*(i) \, \hat{b}_{\alpha}, \qquad \hat{c}_i^{\dagger} = \sum_{\alpha=1}^L \phi_{\alpha}(i) \, \hat{b}_{\alpha}^{\dagger}. \tag{3.63}$$

The unitarity of this transformation ensures the anticommutation relations $\{\hat{b}_{\alpha}, \hat{b}^{\dagger}_{\beta}\} = \delta_{\alpha\beta}$. In terms of these operators the Hamiltonian is diagonal:

$$\hat{\mathcal{H}} = \sum_{\alpha=1}^{L} \epsilon_{\alpha} \, \hat{b}_{\alpha}^{\dagger} \hat{b}_{\alpha} = \sum_{\alpha=1}^{L} \epsilon_{\alpha} \, \hat{n}_{\alpha}, \tag{3.64}$$

where \hat{b}^\dagger_{α} creates a fermion in the α -th Anderson orbital and $\hat{n}_{\alpha}=\hat{b}^\dagger_{\alpha}\hat{b}_{\alpha}$ counts its occupation.

A many-body eigenstate is obtained by choosing a set of orbitals to occupy and acting with the corresponding creation operators on the vacuum. For example, at half filling (L/2 fermions), one occupies L/2 different orbitals. These many-body eigenstates are Slater determinants built from the occupied single-particle eigenstates.

In the site basis, the local number operators transform as

$$\hat{n}_i = \hat{c}_i^{\dagger} \hat{c}_i = \sum_{\alpha,\beta} \phi_{\alpha}(i) \,\phi_{\beta}^*(i) \,\hat{b}_{\alpha}^{\dagger} \hat{b}_{\beta} = \sum_{\alpha,\beta} U_{\alpha\beta}^{(i)} \,\hat{b}_{\alpha}^{\dagger} \hat{b}_{\beta}, \tag{3.65}$$

where $U_{\alpha\beta}^{(i)} \equiv \phi_{\alpha}(i) \, \phi_{\beta}^{*}(i)$. The diagonal terms $\hat{n}_{\alpha} = \hat{b}_{\alpha}^{\dagger} \hat{b}_{\alpha}$ can be expressed via the inverse transformation

$$\hat{n}_{\alpha} = \sum_{i,j} \phi_{\alpha}^{*}(i) \,\phi_{\alpha}(j) \,\,\hat{c}_{i}^{\dagger} \hat{c}_{j} = \sum_{i,j} \tilde{U}_{ij}^{(\alpha)} \,\,\hat{c}_{i}^{\dagger} \hat{c}_{j}, \tag{3.66}$$

with $\tilde{U}_{ij}^{(\alpha)} = \phi_{\alpha}^*(i) \, \phi_{\alpha}(j)$.

It is evident that the site-local operator \hat{n}_i has support in many Anderson orbitals (due to the sum over α, β). If an eigenstate is extended, any local perturbation produces a state with support over many orbitals. For example, creating a particle on an empty site j via $\hat{c}_j^{\dagger}|\phi_{\alpha}\rangle$ can be seen as

$$\hat{c}_{j}^{\dagger}|\phi_{\alpha}\rangle = \sum_{\beta=1}^{L} \phi_{\beta}(j) \,\hat{b}_{\beta}^{\dagger}|\phi_{\alpha}\rangle,\tag{3.67}$$

and for extended states, $\phi_{\beta}(j)$ is generically nonzero for many β , so the excitation mixes many eigenstates.

For localized eigenstates, the situation is different: if $|\phi_{\alpha}\rangle$ is localized around a center R_{α} with localization length $\xi_{\rm loc}$, and if $|R_{\alpha}-j|\gg \xi_{\rm loc}$, then $\phi_{\alpha}(j)$ is exponentially small. A local excitation at j therefore has negligible overlap with $|\phi_{\alpha}\rangle$ and remains confined near j. Formally, the kernel $\tilde{U}_{ij}^{(\alpha)}$ in Eq. (3.66) decays exponentially with $|i-R_{\alpha}|$ and $|j-R_{\alpha}|$, so \hat{n}_{α} is quasilocal in real space—it has appreciable support only on sites within $O(\xi_{\rm loc})$ of R_{α} . In the non-interacting case, these \hat{n}_{α} are exact integrals of motion:

$$[\hat{n}_{\alpha}, \hat{\mathcal{H}}] = 0, \qquad [\hat{n}_{\alpha}, \hat{n}_{\beta}] = 0.$$
 (3.68)

We can then *truncate* the sum over (i, j) in Eq. (3.66) to sites lying within a spatial region centered at R_{α} whose radius is $m \xi_{loc}$, where ξ_{loc} is the localization length of the eigenstate

and m>0 is a dimensionless cutoff parameter controlling how many localization lengths are included. The resulting quasilocal operator $\hat{n}_{\alpha}^{(m)}$ differs from the exact \hat{n}_{α} only by terms supported outside this region; these contributions are exponentially small in m, so the commutator $[\hat{n}_{\alpha}^{(m)}, \hat{\mathcal{H}}]$ vanishes exponentially as $m \to \infty$, recovering the exact conserved quantity.

This property becomes crucial when interactions are introduced:

$$\hat{\mathcal{H}}_{MBL} = -J \sum_{i=1}^{L} \left(\hat{c}_{i}^{\dagger} \hat{c}_{i+1} + \text{h.c.} \right) + \sum_{i=1}^{L} \varepsilon_{i} \, \hat{n}_{i} + \Delta \sum_{i=1}^{L} \hat{n}_{i} \hat{n}_{i+1} , \qquad (3.69)$$

for small Δ , the exact integrals of motion of the interacting Hamiltonian are expected to be weakly 'dressed' versions of these non-interacting \hat{n}_{α} , and their quasilocality ensures that such dressed operators can remain spatially localized and mutually commuting in the MBL phase [91–94, 239–242]. Concretely, we can regard the quasilocal integrals of motion (LIOMs) in the MBL phase as

$$\hat{I}_{\alpha} = q \,\hat{n}_{\alpha} + \sum_{m=1}^{\infty} U_{\alpha}^{(m)} \,\hat{O}_{\alpha}^{(m)} \,, \tag{3.70}$$

where q is the finite overlap between the interacting LIOM \hat{I}_{α} and its non-interacting counterpart $\hat{n}_{\alpha} = \hat{b}_{\alpha}^{\dagger}\hat{b}_{\alpha}$. The operators $\hat{O}_{\alpha}^{(m)}$ contain up to (2m+1)-body terms whose support is centered around the localization center R_{α} of orbital α and extends over a region of radius $m\,\xi_{\rm op}$. The coefficients $U_{\alpha}^{(m)}$ decay exponentially with $m,\,U_{\alpha}^{(m)}\sim e^{-m/\xi_{\rm op}}$, reflecting the quasilocal nature of \hat{I}_{α} .

Notice that $\xi_{\rm op}$ is the localization length associated with the LIOM operators, which generally differs from the single-particle localization length $\xi_{\rm loc}$. The two coincide, $\xi_{\rm op}=\xi_{\rm loc}$, only in the non-interacting limit $\Delta=0$. For $\Delta>0$, interactions cause the LIOMs to acquire multi-site and multi-body terms through dressing, which increases their spatial extent and generally yields $\xi_{\rm op}>\xi_{\rm loc}$.

These LIOMs share a common eigenbasis with the Hamiltonian and thus form a complete set of commuting operators in terms of which the Hamiltonian can be expanded. In the presence of interactions, the Hamiltonian is no longer quadratic in fermionic operators, and its most general diagonal form in the LIOM basis reads

$$\hat{\mathcal{H}}_{\text{MBL}} = h_0 \,\hat{\mathbb{I}} + \sum_{\alpha} h_{\alpha} \,\hat{I}_{\alpha} + \sum_{\alpha,\beta} h_{\alpha\beta} \,\hat{I}_{\alpha} \hat{I}_{\beta} + \sum_{\alpha,\beta,\gamma} h_{\alpha\beta\gamma} \,\hat{I}_{\alpha} \hat{I}_{\beta} \hat{I}_{\gamma} + \cdots , \qquad (3.71)$$

where the sums run over all LIOM indices and the expansion includes interactions of arbitrarily high order between them.

Phenomenological models for the MBL phase [239, 240] aim to capture its essential static and dynamical properties without relying on a full microscopic solution. A common approach is to represent the system in terms of this complete set of quasilocal integrals of motion \hat{I}_{α} , and assume that they possess a binary spectrum and can be interpreted as effective spins or occupation number operators (so called ℓ -bits). The full local operator algebra can then be generated from these \hat{I}_{α} via suitable ladder operators [275].

In this picture, many-body eigenstates are simultaneous eigenstates of all \hat{I}_{α} , and the quasilocality of the latter ensures spatial localization, absence of transport, and persistence of local information at arbitrarily long times. When the system enters the delocalized phase, the \hat{I}_{α} cease to be quasilocal and instead become extended over the entire system, losing their role as a complete set of local conserved quantities.

3.3 The Random Field XXZ Chain

In the present work, we will consider the XXZ spin- $\frac{1}{2}$ chain with interaction anisotropy parameter Δ ($0 \le \Delta \le 1^2$), and a random magnetic field h_i uniformly drawn from [-W, W]. The Hamiltonian, for a chain of L sites, is given by

$$\hat{\mathcal{H}} = \sum_{i=1}^{L} \left(\hat{S}_{i}^{x} \hat{S}_{i+1}^{x} + \hat{S}_{i}^{y} \hat{S}_{i+1}^{y} + \Delta \hat{S}_{i}^{z} \hat{S}_{i+1}^{z} + h_{i} \hat{S}_{i}^{z} \right) . \tag{3.72}$$

For $\Delta=1$, we recover the strongly interacting random-field Heisenberg chain, which serves as the paradigmatic model of the MBL transition. This model is equivalent, through a Jordan-Wigner transformation

$$\hat{S}_{j}^{x} = \frac{1}{2} \left(\hat{c}_{j}^{\dagger} + \hat{c}_{j} \right) \prod_{k=1}^{j-1} \left(1 - 2 \, \hat{c}_{k}^{\dagger} \hat{c}_{k} \right), \tag{3.73}$$

$$\hat{S}_{j}^{y} = -\frac{i}{2}(\hat{c}_{j}^{\dagger} - \hat{c}_{j}) \prod_{k=1}^{j-1} \left(1 - 2\,\hat{c}_{k}^{\dagger}\hat{c}_{k}\right), \tag{3.74}$$

$$\hat{S}_{j}^{z} = \hat{c}_{j}^{\dagger} \hat{c}_{j} - \frac{1}{2} \hat{\mathbb{I}} , \qquad (3.75)$$

to that of interacting spinless fermions hopping on a chain of L sites, written as

$$\hat{\mathcal{H}} = \frac{1}{2} \sum_{i=1}^{L} \left(\hat{c}_i^{\dagger} \hat{c}_{i+1} + \text{h.c.} + 2\Delta \hat{n}_i \hat{n}_{i+1} + 2h_i \hat{n}_i \right) , \qquad (3.76)$$

with $\hat{c}_i^{\dagger}(\hat{c}_i)$ the creation (annihilation) fermionic operators and $\hat{n}_i = \hat{c}_i^{\dagger}\hat{c}_i$ the local occupation number operator.

In both cases, setting $\Delta=0$ corresponds to the non-interacting limit, which—when described in the language of spinless fermions—is known to exhibit Anderson localization for any nonzero value of the disorder strength W. In this regime, the localization length of the single-particle Anderson-localized orbitals diverges as $1/W^2$ in the limit $W\to 0$. The Hamiltonian preserves the U(1) symmetry thus conserving the magnetization in the z direction—or total particle number—of the system. In the present work, we restrict to studying the sector of zero magnetization, or half-filling, with periodic boundary conditions.

3.3.1 The Hilbert space picture

The Hilbert space structure of the model described above has been extensively studied [90, 283, 349–353]. Any quantum state $|\Psi\rangle$ can be expressed as superpositions of the many-body basis states in the form

$$|\Psi\rangle = \sum_{I=1}^{\mathcal{N}} \psi_I |I\rangle ,$$
 (3.77)

where $\{|I\rangle\}$ is the appropriate set of many-body basis states. We will focus on two bases: spin configurations $\{|I\rangle\}_S$ and the Anderson basis $\{|I\rangle\}_A$. These are the most relevant for the model in Eq. (3.72), as each diagonalizes the Hamiltonian in specific limits where

²The sign of the interaction is not relevant at high energy, see for instance Ref. [348]

localization is well-understood, the infinite-disorder limit $W \to \infty$ for $\{|I\rangle\}_S$, and the non-interacting limit $\Delta = 0$ for $\{|I\rangle\}_A$.

Once the computational basis is chosen, the Hamiltonian (3.72)-(3.76) can be recast as

$$\hat{\mathcal{H}} = \sum_{I} E_{I} |I\rangle\langle I| + \sum_{\langle IJ\rangle} T_{IJ} |I\rangle\langle J| . \tag{3.78}$$

This is now a tight-binding model of a fictitious single-particle hopping on a graph. The number of vertices $\mathcal N$ in this high-dimensional graph represents the possible basis states associated to the problem (i.e., the Hilbert space volume). In a system with L spins or fermions, the dimension of Hilbert space grows exponentially with L. In the zero magnetization (or half-filling) sector it is given by $\mathcal N = \binom{L}{L/2} \approx 2^L \sqrt{\frac{2}{\pi L}}$. Each of these states have an associated 'on-site' energy E_I , given by the expectation value of the diagonal part of the Hamiltonian in the chosen basis, which is random through the fields h_i . The edges of the Hilbert space graph joining different vertices are given by the non-zero off-diagonal elements T_{IJ} , which gives the 'hopping' amplitudes of the fictitious particle between two neighboring vertices. A pictorial representation of the Hilbert space graph of the XXZ model—considering spin configurations $|I\rangle_S$ —with L=8 is shown in Fig. 3.2.

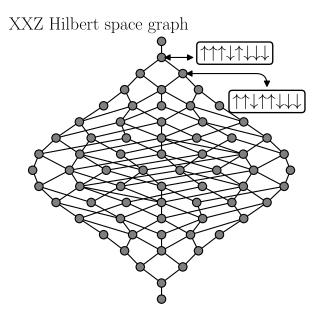


Fig. 3.2: Sketch of Hilbert space graph for the XXZ model with L=8 spins. Adapted from Ref. [353].

Each vertex—basis state—is connected to a large number of other vertices via local interactions, such as spin flips or particle hoppings. Even though the underlying physical system is one-dimensional, the connectivity in Hilbert space is governed by the action of local operators on basis states. This gives rise to a complex, high-dimensional network in which each vertex (or basis state) is connected to a number of others that scales with L or a power of L, depending on the chosen computational basis (see below).

This reformulation provides a framework in which spectral and dynamical properties of the many-body quantum disordered system can be understood in terms of the effective single-particle problem on an underlying complex graph structure. Generically, this Hilbert space graph exhibits an effective infinite dimensionality: the number of vertices at a given distance from a reference vertex grows exponentially with that distance.

The main differences compared to the Anderson localization problem on high dimensional graphs [352, 354–356] lie in the scaling and structure of the problem. In the many-body case, the corresponding Hilbert space (or Fock-space) graph is deterministic and contains loops of various lengths. Moreover, the many-body problem features strongly correlated on-site energies E_I and matrix elements T_{IJ} , further complicating the analysis compared to standard single-particle Anderson localization on tree-like structures such as the Bethe lattice.

In the following sections we introduce and explain in detail the two computational bases used to investigate the random-field XXZ model.

3.3.2 Spin basis

The first computational basis we will be working with is the simultaneous eigenstates of the S_i^z operators: $\{|I\rangle\}_S$ are just classical Ising spin configurations on the chain with zero global magnetization in the z. For example, for L=6 one of its basis states is $|\uparrow\downarrow\downarrow\uparrow\uparrow\downarrow\rangle$. The states in the set $\{|I\rangle\}_S$ are the eigenstates of $\hat{\mathcal{H}}_0 = \sum_{i=1}^L (\Delta \hat{S}_i^z \hat{S}_{i+1}^z + h_i \hat{S}_i^z)$ with eigenenergy $E_I = \langle I|\hat{\mathcal{H}}_0|I\rangle$. In this basis the off-diagonal elements, T_{IJ} , are not random and are in fact all set to $T_{IJ} = 1/2$ for those basis states—I and J—that differ by the exchange of two spins forming domain walls in the chain. It is clear that the connectivity of the vertices is not constant and depends on the number of domain walls in each basis state. For example, the Néel state $|\uparrow\downarrow\uparrow\downarrow\dots\uparrow\downarrow\rangle$, and its time-reversed symmetric, are maximally connected as they both have L domain walls, but these two Néel states are distant from each other by L/2 spin flips. Instead, the state $|\uparrow\uparrow\dots\uparrow\downarrow\downarrow\dots\downarrow\rangle$ has only two domain walls (with periodic boundary conditions) thus having just two neighboring vertices in the Hilbert space graph. The average degree of the Hilbert space graph is L/2.

Through a Jordan-Wigner transformation, these basis states can be mapped to Fock states of spinless fermions hopping on a lattice at half-filling. For example, the Néel state introduced above becomes $|1\ 0\ 1\ 0\ \dots\ 1\ 0\rangle$ in this representation. Throughout the text, we will refer to both representations as the *spin basis*, and we will switch between the conventional spin configuration notation (e.g., $|\uparrow\downarrow\uparrow\downarrow\dots\uparrow\downarrow\rangle$) and the corresponding *Fock state* notation (e.g., $|1\ 0\ 1\ 0\ \dots\ 1\ 0\rangle$) whenever it is convenient for clarity.

3.3.3 Anderson basis

Another natural choice that, by construction, captures the fact that the system remains localized at any finite disorder in the absence of interactions is the Anderson basis of single particle localized orbitals, described below (see also Refs. [263, 357–360]). This second computational basis corresponds to the eigenstates of the non-interacting part of the Hamiltonian in Eq. (3.76), defined as $\hat{\mathcal{H}}_{\rm NI} = \frac{1}{2} \sum_{i=1}^L (\hat{c}_i^{\dagger} \hat{c}_{i+1} + \text{h.c.} + 2h_i \hat{n}_i)$. This $\hat{\mathcal{H}}_{\rm NI}$ defines the Anderson model of a spinless fermion hopping on a chain, in a random potential. Eigenpairs of this single-particle Hamiltonian stem from $\hat{\mathcal{H}}_{\rm NI}\phi_{\alpha} = \epsilon_{\alpha}\phi_{\alpha}$, and are known to be Anderson localized. The single-particle orbitals ϕ_{α} are used to construct the unitary transformation that diagonalizes the non-interacting Hamiltonian $\hat{\mathcal{H}}_{\rm NI}$ as it was already discussed in Sec. 3.2.5,

$$\hat{\mathcal{H}}_{NI} = \sum_{\alpha} \epsilon_{\alpha} \hat{b}_{\alpha}^{\dagger} \hat{b}_{\alpha} , \qquad \hat{b}_{\alpha} = \sum_{i=1}^{L} \phi_{\alpha}(i) \hat{c}_{i} . \qquad (3.79)$$

The basis states are thus defined as the occupation numbers of the L single particle orbitals, i.e., the tensor product of the simultaneous eigensates of the number operators $\hat{b}^{\dagger}_{\alpha}\hat{b}_{\alpha}$. The

Anderson basis $\{|I\rangle\}_A$, is thus built as

$$|I\rangle = \prod_{\alpha=1}^{L} (\hat{b}_{\alpha}^{\dagger})^{n_{\alpha}} |\varnothing\rangle , \qquad E_{I}^{\text{NI}} = \sum_{\alpha} n_{\alpha} \epsilon_{\alpha} , \qquad (3.80)$$

where $|\varnothing\rangle$ is the vacuum state, and n_α is the fermion occupation number for the α -th orbital, that in our case for half-filling fulfills $\sum_{\alpha} n_{\alpha} = L/2$. $E_I^{\rm NI}$ is the non-interacting energy associated to the basis state $|I\rangle$. The interacting part, $\sum_{i=1}^L \hat{n}_i \hat{n}_{i+1}$, transforms under the unitary transformation defined in Eq. (3.79) into

$$\hat{V} = \sum_{\alpha\beta\gamma\delta} V_{\alpha\beta\gamma\delta}(\Delta) \; \hat{b}_{\alpha}^{\dagger} \hat{b}_{\beta}^{\dagger} \hat{b}_{\gamma} \hat{b}_{\delta}$$
 (3.81)

with

$$V_{\alpha\beta\gamma\delta}(\Delta) = \Delta \sum_{i=1}^{L} \phi_{\alpha}^{*}(i)\phi_{\beta}^{*}(i+1)\phi_{\gamma}(i+1)\phi_{\delta}(i) . \tag{3.82}$$

We can unfold Eq. (3.82) to consider explicitly the three non-vanishing contributions to $V_{IJ} = \langle J | \hat{V} | I \rangle$, according to the groupings of indices α, β, γ and δ : (i) $\alpha = \delta$ and $\beta = \gamma$, gives a diagonal contribution in the form

$$\hat{V}_d = 2\sum_{\alpha>\beta} \left(V_{\alpha\beta\beta\alpha} - V_{\beta\alpha\alpha\beta}\right) \hat{n}_{\beta} \hat{n}_{\alpha} , \qquad (3.83)$$

making the on-site energy of the associated basis state $|I\rangle$ to be $E_I = E_I^{\rm NI} + \langle I|\hat{V}_d|I\rangle$. The other two non-zero contributions to V_{IJ} come from the off-diagonal entries $|I\rangle \neq |J\rangle$, that will construct the hopping terms T_{IJ} . The second contribution comes from (ii) $\alpha = \delta, \beta \neq \gamma$:

$$\hat{V}_1 = 2\sum_{\alpha\beta\gamma} (V_{\alpha\beta\gamma\alpha} - V_{\beta\alpha\alpha\gamma}) \hat{n}_{\alpha} \hat{b}_{\beta}^{\dagger} \hat{b}_{\gamma} , \qquad (3.84)$$

where the occupation of the γ -th orbital has been replaced by a new occupation in the orbital labeled by β , given that $n_{\alpha}=n_{\gamma}=1$ and $n_{\beta}=0$. This 'assisted' hopping connects each vertex with $L^2/4$ nearest-neighbors. The third and final contribution corresponds to (iii) $\alpha \neq \beta \neq \gamma \neq \delta$:

$$\hat{V}_{2} = \sum_{\substack{\alpha > \beta \\ \gamma > \delta}} (V_{\alpha\beta\gamma\delta} + V_{\beta\alpha\delta\gamma} - V_{\alpha\beta\delta\gamma} - V_{\beta\alpha\gamma\delta}) \hat{b}_{\alpha}^{\dagger} \hat{b}_{\beta}^{\dagger} \hat{b}_{\gamma} \hat{b}_{\delta} . \tag{3.85}$$

This contribution is non-zero provided that $n_{\delta}=n_{\gamma}=1$ and $n_{\alpha}=n_{\beta}=0$. This contribution adds $L^4/64-L^3/16+L^2/16$ neighboring vertices to each single vertex. Thus, the Fock-space graph has constant connectivity $z=L^2/4$ [$(L/2-1)^2/4+1$], for each basis state. Note that unlike the spin basis, here the T_{IJ} are all random, and broadly distributed [357].

3.4 Standard observables and their limitations

The characterization of the many-body localization (MBL) transition has traditionally relied on a set of widely studied observables, including imbalance dynamics, autocorrelation functions, adjacent gap ratios, and entanglement scaling laws. These diagnostics have provided essential evidence for ergodicity breaking and are routinely employed in both numerical [102, 246, 247] and experimental studies [361]. Yet their interpretation remains subtle: signatures that appear robust at finite sizes can become ambiguous when probed at longer times or larger system sizes.

These ambiguities, together with discrepancies among different diagnostics, have motivated a better understanding of rare Griffiths regions—spatial domains where the disorder is anomalously weak [100, 243, 248–254]. These regions can act as ergodic seeds that grow with system size, potentially destabilizing localization in the thermodynamic limit [95–97, 100–103, 243–245].

In the following, we review standard probes of the MBL transition and introduce the two primary frameworks that describe how rare Griffiths regions can destabilize the MBL phase: many-body resonances and the avalanche mechanism.

Spectral statistics

Within the ETH phase, the eigenstates/values of the Hamiltonian follow the predictions of random matrix theory. For example, as the matrix elements of Eq. (3.72) are real, its eigenvalues should follow the Gaussian Orthogonal Ensemble (GOE) statistics. As a result, its eigenvalues (E_{α}) follow a Wigner semi-circle law while the level spacings ($\delta E_{\alpha} = E_{\alpha} - E_{\alpha-1}$) are distributed according to the Wigner-Dyson law.

However, in many-body systems, the energy levels are not uniformly spaced—the density of states varies with energy. As a result, the average level spacing $\langle \delta E_{\alpha} \rangle$ may introduce finite-size effects and requires arbitrary choices—as selecting a window size for the local averaging. Instead, is better to study the ratio of consecutive level spacings r_n , introduced by Huse and Oganesyan in Ref. [362]

$$r_{\alpha} = \frac{\min(\delta E_{\alpha}, \delta E_{\alpha-1})}{\max(\delta E_{\alpha}, \delta E_{\alpha-1})}.$$
(3.86)

Whose average in the GOE ensemble is $\langle r \rangle_{\rm GOE} \simeq 0.5307$. In the MBL phase the eigenvalues are uncorrelated, and thus they follow a Poisson distribution, for which one obtains $\langle r \rangle_{\rm Poisson} \simeq 0.386$. Where in the case of MBL systems the average $\langle r \rangle$ is computed over disorder realizations and different eigenstates within the same energy. This measure has been extensively used as an indicator of the absence of level repulsion, generically regarded as a signature of integrability [363–367]. We show an example of this calculation for the Hamiltonian in Eq. (3.72), adapted from the work of Luitz, Laflorencie, and Alet [246].

An argument challenging the conclusions drawn from this observable was presented in Ref. [95]. There, it was shown that close to the ETH regime—i.e., for $\langle r \rangle$ $(W) \gtrsim \langle r \rangle_{\text{GOE}}$ —the curves collapse when rescaled as $\langle r \rangle$ $(W) \to \langle r \rangle$ (W/L). This scaling implies that $\langle r \rangle$ remains essentially constant if the disorder strength W is increased proportionally to the system size L. Consequently, in the thermodynamic limit $(L \to \infty)$ one would have $\langle r \rangle$ $(W) \to \langle r \rangle_{\text{GOE}}$ for any finite W, effectively reducing the MBL phase to a finite-size crossover rather than a genuine phase transition.

Entanglement entropy

Consider an isolated quantum system S that can be partitioned into two subsystems, A and B ($S = A \cup B$). The central implication of the ETH is that if subsystem A is sufficiently small, its complement B effectively acts as a thermal bath for A. This principle directly influences the structure of the eigenstates $|\psi_{\alpha}\rangle$ of the system: for any local observable within A, the expectation values $\langle \psi_{\alpha} | \hat{\mathcal{O}}_A | \psi_{\alpha} \rangle$ are thermal. Consequently, the von Neumann entanglement entropy $S_A^{(\alpha)}$, obtained by tracing out the degrees of freedom of subsystem B,

$$S_A^{(\alpha)} = -\text{Tr}(\hat{\rho}_A \ln \hat{\rho}_A) , \quad \text{with} \quad \hat{\rho}_A \equiv \text{Tr}_B |\psi_\alpha\rangle \langle\psi_\alpha| ,$$
 (3.87)

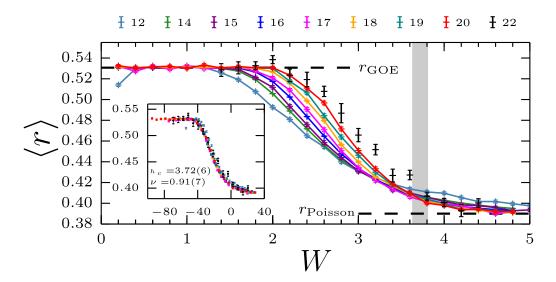


Fig. 3.3: Adjacent gap ratio as a function of disorder strength W in the middle of the spectrum. Inset: data collapse used to extract the critical disorder strength W_c and exponent ν . Figure adapted from Ref. [246]

needs to be extensive with respect to the subsystem size L_A (in a spin chain of size L then $L = L_A + L_B$, for example). This behavior is known as volume-law scaling: for a system of spatial dimension d and linear size L, the entanglement entropy scales proportionally to the system's volume, i.e., as L^d . Such volume-law scaling has been confirmed in disorder-free quantum spin chains [368–372].

However, Bauer and Nayak observed in Ref. [373] that within the MBL phase, the finite localization length of the eigenstates causes a breakdown of this volume-law behavior. Instead, the entanglement entropy follows an area-law scaling, growing proportionally to the surface area of the subsystem—corresponding to a constant in one-dimensional systems. This property has become one of the standard diagnostics for identifying the MBL transition [246, 270, 374–377].

To make meaningful comparisons of entanglement entropy at different system sizes, it is convenient to rescale the entanglement entropy by the value predicted from random matrix theory $S_{\rm RMT}$, in the respective subspace of the Hilbertspace, i.e. $s=S_A/S_{\rm RMT}$. The average over disorder realizations and mid-spectrum energies $\langle s \rangle$ may serve as an indicator of the transition.

This was studied for the $J_1 - J_2$ model, another spin chain that exhibits an ETH-MBL transition [246] with Hamiltonian

$$\hat{\mathcal{H}}_{J_1 - J_2} = \hat{\mathcal{H}}_{XXZ} + \sum_{i=1}^{L} \left(\hat{S}_i^x \hat{S}_{i+2}^x + \hat{S}_i^y \hat{S}_{i+2}^y + \Delta \hat{S}_i^z \hat{S}_{i+2}^z \right)$$
(3.88)

where $\hat{\mathcal{H}}_{XXZ}$ is the original model defined in Eq. (3.72). In the MBL phase we expect $\langle s \rangle \to 0$, whereas for ETH $\langle s \rangle \to 1$, upon increasing L, as show in ??.

Notably, the crossing point of the $\langle s \rangle (W)$ curves exhibits a pronounced drift as the system size increases. This behavior is analogous to that observed for the averaged gap ratio $\langle r \rangle$ discussed earlier, preventing unambiguous conclusions about the fate of the crossover in the thermodynamic limit. Furthermore, it has been reported [94] that the magnitude of this drift can depend on the specific observable considered.

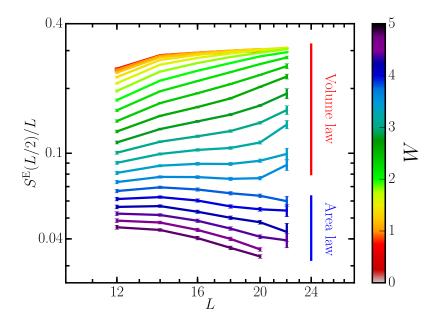


Fig. 3.4: Exact diagonalization results for the disorder-average entanglement entropy computed at half-chains using the partition of the subsystem as half the chaine L/2 for the Hamiltonian in Eq. (3.72) with $\Delta=1$ at a fixed energy density in the middle of the many-body spectrum. One sees a qualitative change from volume-law at small disorder W to an area law at stronger disorders. Figure adapted from Ref. [246]

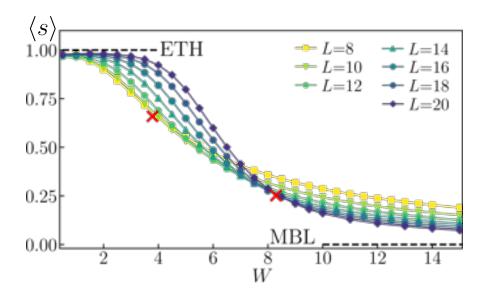


Fig. 3.5: Rescaled entanglement $\langle s \rangle$ of mid-spectrum eigenstates of the J_1 – J_2 model at the ETH-MBL crossover. Figure taken from Ref. [94].

Some dynamical properties

The breakdown of ergodicity of MBL is reflected directly in its dynamics, and there are many observables that attempt to measure this directly. The most intuitive quantity is the autocorrelation function—often used in the study of glassy systems. Similarly one can define the autocorrelation as

$$C_0(t) = \frac{1}{L} \operatorname{Re} \sum_{i=1}^{L} \operatorname{Tr} \left(\hat{\rho}_0 \, \hat{S}_i^z(t) \hat{S}_i^z(0) \right) , \qquad (3.89)$$

where the real part is taken to ensure the observed quantity is hermitian, and $\hat{\rho}_0$ is the associated density matrix $e^{-\beta \hat{\mathcal{H}}}/\mathcal{Z}$.

The infinite time average will then be given by

$$\lim_{T \to \infty} \int_0^T C_0(t')dt' = \frac{1}{L} \sum_i \sum_{\alpha} e^{-\beta E_{\alpha}} |\langle \psi_{\alpha} | \hat{S}_i^z | \psi_{\alpha} \rangle|^2 = q_{\text{EA}} , \qquad (3.90)$$

where $|\psi_{\alpha}\rangle$ and E_{α} are the eigenvalue and eigenstate of the Hamiltonian, and $q_{\rm EA}$ is the Edwards-Anderson (EA) order parameter. Equivalently as in spin-glasses, if we in the ergodic phase (where ETH is obeyed) we expect $q_{\rm EA}=0$, as the expected value $\langle \psi_{\alpha}|\hat{S}_{i}^{z}|\psi_{\alpha}\rangle$ will match the thermal expectation for every eigenstate. In the MBL phase instead, localized eigenstates will produce a saturation of the EA parameter to a non-zero value. For systems with no spin reflection symmetry, one could use this as an order parameter of the MBL phase [374, 378–381].

Within this context, one can examine the autocorrelation function setting as initial state an eigenstate of the Hamiltonian. In this case, the density matriz is just $\hat{\rho}_0 = |\psi_{\alpha}\rangle \langle \psi_{\alpha}|$ and we obtain

$$C_0(t) = \frac{1}{L} \sum_{i=1}^{L} \langle \psi_{\alpha} | \hat{S}_i^z(t) \hat{S}_i^z(0) | \psi_{\alpha} \rangle .$$
 (3.91)

We can split the operator \hat{S}_i^z into its static value and its fluctuations

$$\hat{S}_{i}^{z} = \langle \psi_{\alpha} | \hat{S}_{i}^{z} | \psi_{\alpha} \rangle + \delta \hat{S}_{i}^{z} , \qquad \text{with} \quad \delta \hat{S}_{i}^{z} = \hat{S}_{i}^{z} - \langle \psi_{\alpha} | \hat{S}_{i}^{z} | \psi_{\alpha} \rangle , \qquad (3.92)$$

making the time dependence of operators in Eq. (3.91) to be entirely contained in their respective fluctuations $\delta \hat{S}_i^z(t)$. As a result we get

$$\langle \psi_{\alpha} | \hat{S}_{i}^{z}(t) \hat{S}_{i}^{z}(0) \psi_{\alpha} \rangle = \langle \psi_{\alpha} | \hat{S}_{i}^{z} | \psi_{\alpha} \rangle^{2} + \langle \psi_{\alpha} | \delta \hat{S}_{i}^{z}(t) \delta \hat{S}_{i}^{z}(0) | \psi_{\alpha} \rangle . \tag{3.93}$$

After time averaging, the EA order parameter is then

$$q_{EA} = \frac{1}{L} \sum_{i=1}^{L} \langle \psi_{\alpha} | \hat{S}_i^z | \psi_{\alpha} \rangle^2 . \qquad (3.94)$$

It is then interesting to analyze the time evolution of the fluctuating state $|\tilde{\psi}_{\alpha}\rangle = \delta \hat{S}_i^z |\psi_{\alpha}\rangle$, specifically its return probability

$$\mathscr{R}(t) = |\langle \tilde{\psi}_{\alpha} | \delta e^{-i\hat{\mathcal{H}}t} | \tilde{\psi}_{\alpha} \rangle|^2 = |\langle \tilde{\psi}_{\alpha} | \hat{S}_i^z \delta e^{-i\hat{\mathcal{H}}t} \delta \hat{S}_i^z | \psi_{\alpha} \rangle|^2. \tag{3.95}$$

This scenario is equivalent to the situation explained in Sec. 3.2.5, in which a local excitation (\hat{S}_i^z) changes the eigenstate $|\psi_\alpha\rangle$. The return probability then quantifies the spread of this local excitation: in the ETH phase the excitation spreads over the entire system, the initial fluctuating state is forgotten and the return probability goes to zero for asymptomatically large times. In the MBL phase, the local excitation connects the perturbed eigenstate $|\tilde{\psi}_\alpha\rangle$ only to a few nearby eigenstates (those compatible with the local integrals of motion). As the fluctuating state remains confined near the site i, in asymptotically large times we get a finite return probability.

In the Hilbert (Fock) space framework, several attempts have been made to extend the characterization of single-particle eigenstate statistics—described via the inverse participation ratio (IPR), see Sec. 3.1.2—to many-body eigenstates. In this case, the size of real space is replaced by the Hilbert space volume $\mathcal N$

$$\mathcal{I}_{q}^{(\alpha)} = \sum_{I} |\langle \psi_{\alpha} | I \rangle|^{2q}$$
 (3.96)

where $|I\rangle$ is some chosen basis state. This quantity is closely related to the basis-dependent Réyni participation entropies [246, 382, 383].

An attempt to study this quantity (for q=2) in the MBL problem was performed in Ref. [384], using the basis state $\hat{S}^z_i | \psi_\alpha \rangle$. Consequently, we can relate this IPR to the return probability. By expanding $\tilde{\psi}_\alpha$ over the relevant eigenbasis $|\psi_\beta\rangle$ and taking the infinite time average of the return probability we get

$$\mathcal{I}_{2}^{(\alpha)} = \lim_{T \to \infty} \int_{0}^{T} dt \left| \sum_{\beta} |M_{\alpha\beta}|^{2} e^{-iE_{\beta}t} \right|^{2} = \sum_{\beta} |M_{\alpha\beta}|^{4} , \qquad (3.97)$$

with $M_{\alpha\beta} = \langle \psi_{\alpha} | \delta \hat{S}_i^z | \psi_{\beta} \rangle$.

The inverse paticipation ratio scales as $\mathcal{I}_2^{(\alpha)} \propto \mathcal{N}^{-\tilde{D}_2}$, where \tilde{D}_2 is a generalized dimension [260, 384]. Analogous to the single particle case, as $L \to \infty$ the IPR $\mathcal{I}_2 \to 0$ in the ergodic phase, whereas \mathcal{I}_2 remains finite in the localized phase, consistent with the criterion coming from the return probability.

Several other dynamical probes can be considered, including conductivities, mean-square displacements, and dynamical structure factors—the latter being widely employed in the study of structural glasses [19, 21]. A particularly relevant observable in this context is the imbalance, defined as

$$\mathscr{I}(t) = \sum_{i=1}^{L} (-1)^i \langle \hat{S}_i^z(t) \rangle;, \qquad (3.98)$$

which can be directly measured in cold-atom experiments [361]. In these setups, site occupations are accessed via band-mapping techniques, enabling the study of post-quench dynamics of a charge-density wave initialized in a Néel state (i.e., $|\uparrow\downarrow\uparrow\downarrow...\rangle$). This quantity has also been employed in numerical investigations [385].

It has been pointed out in the literature [94, 102] that a finite long-time value of the autocorrelation function does not necessarily signal the presence of a stable MBL phase. In systems with very slow dynamics, finite-size effects can lead to a saturation of $C_0(t)$ over the time scales accessible in numerical simulations or experiments, even if the system is ultimately thermalizing. In such cases, the decay of correlations may occur only on times far beyond those probed, producing apparent plateaus—such as those visible in Fig. 3.6—that mimic the behavior expected from truly localized states. Similar considerations apply to other commonly used diagnostics, such as the return probability and the imbalance, whose finite-time persistence may also reflect transient rather than asymptotic localization.

Entanglement growth

If a many-body system is initialized in a state with low entanglement—for example a simple basis state $|I\rangle = |\uparrow\downarrow\downarrow\uparrow\downarrow\ldots\rangle$ —the unitary evolution will typically generate entanglement between increasingly distant regions as time evolves.

Heuristically, one can understand this studying the time evolution of a simple system with two spins. We prepare the system to be initially in a basis state $|\uparrow_A\rangle\otimes|\downarrow_B\rangle$. This state is untangled: the reduced density matrix is pure $\hat{\rho}_A=|\uparrow\rangle\langle\uparrow|$ (or equivalent $\hat{\rho}_B=|\downarrow\rangle\langle\downarrow|$). Now assume that the Hamiltonian has eigenstates

$$|\psi_{1}\rangle = \frac{1}{\sqrt{2}}(|\uparrow\downarrow\rangle + |\downarrow\uparrow\rangle) ,$$

$$|\psi_{2}\rangle = \frac{1}{\sqrt{2}}(|\uparrow\downarrow\rangle - |\downarrow\uparrow\rangle) ,$$
(3.99)

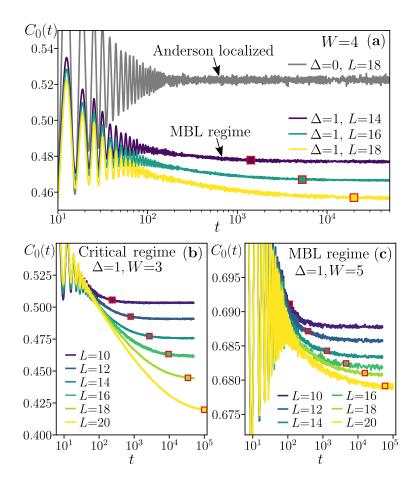


Fig. 3.6: Time evolution of the density autocorrelation function C(t) in the disordered XXZ spin-1/2 chain (3.72). (a) For W =4, the interacting case ($\Delta =$ 1) exhibits a slow decay of C(t) up to the Heisenberg time t_H (red squares), while for $\Delta = 0$ (Anderson localized phase) C(t) saturates after initial oscillations. (b) At W=3, a pronounced decay is seen, consistent with the crossover regime. (c) At W=5 (MBL regime, $W > W^*(L)$ for $12 \leq L \leq 20$), C(t) decays extremely slowly; whether it ultimately vanishes in the $L, t \rightarrow \infty$ limit remains unresolved. Adapted from Refs. [94, 102].

with eigenenergies E_1 and E_2 , respectively. The initial basis state is a superposition of these two eigenstates $|I\rangle=\frac{1}{\sqrt{2}}(|\psi_1\rangle+|\psi_2\rangle)$. Hence, applying the time evolution operator we obtain

$$|\Psi(t)\rangle = \hat{\mathcal{U}}(t) |\uparrow\downarrow\rangle = e^{-i\hat{\mathcal{H}}t} |\uparrow\downarrow\rangle$$

$$= \frac{1}{2} \left[e^{-iE_1t} (|\uparrow\downarrow\rangle + |\downarrow\uparrow\rangle) + e^{-iE_2t} (|\uparrow\downarrow\rangle - |\downarrow\uparrow\rangle) \right]$$

$$= \frac{1}{2} \left[(e^{-iE_1t} + e^{-iE_2t}) |\uparrow\downarrow\rangle + (e^{-iE_1t} - e^{-iE_2t}) |\downarrow\uparrow\rangle \right] . \tag{3.100}$$

We can see from Eq. (3.100) that at later times the relative phases $e^{-i(E_1\pm E_2)t}$ become non-trivial, and the states becomes a coherent superposition of entangled basis states. This produces the reduced density matrix to become 'mixed' and the entanglement entropy grows.

This irreversible growth of entanglement—quantified by the growth of the von Neumman entropy in Eq. (3.87)—is an essential feature of thermalization, and as a result has been addressed in diverse contexts ranging from conformal field theory [386–388] and holography [389–394] to integrable [395–397], nonintegrable [398, 399] and, our main interest, strongly disordered spin chains [240, 385, 400–402].

In non-interacting systems [386], the entanglement entropy S(t), as defined in Eq. (3.87), increases linearly with time. This defines a ballistic growth of the entanglement entropy, also observed in integrable systems—such as the clean XXZ chain—after a global quench [386, 403] and non-integrable systems [399]. The introduction of disorder renders this growth slower, as first observed in Ref. [403] for an XX spin chain with random bonds.

In the MBL phase, distant spins at a distance n are weakly coupled and they get slowly entangled at long times, this leads to a very slow entanglement growth of $S(t) \sim \ln t$ [400,

401] making it consistent with the framework of local integrals of motion [239]. The crossover between ETH and MBL phases exhibits an entanglement entropy that grows sub-ballistically, with $S(t) \sim t^{1/z(W)}$ with a disorder dependent exponent [385].

Nevertheless, a slight systematic increase in the exponent z can be observed as the time and length scales grow [404, 405]. This observation raises the question of whether a gradual onset of slow thermalization is taking place, wherein the logarithmic growth of S(t) is replaced by a slow algebraic increase of the entanglement entropy. Distinguishing between algebraic and logarithmic growth is highly non-trivial when the dynamics become slower with increasing W, particularly since the available data typically cover only a limited time window.

3.4.1 The destabilization of the MBL phase

The aforementioned observables of the MBL phase suffer from significant limitations and strong finite-size effects, which have prevented studies from reaching definitive and unambiguous conclusions about the fate of the MBL phase in the thermodynamic limit. Furthermore, mechanisms that restore ergodicity within the MBL phase have been proposed; these are inherently non-perturbative in nature. As a result, both perturbative descriptions of the MBL transition and these standard observables have been met with skepticism. These non-perturbative effects are associated to the existence of regions of the chain that are effectively ergodic, for which the disorder is anomalously weak. Such Griffiths regions are expected to necessarily occur in sufficiently large systems, leading to the idea that they can trigger mechanisms that restore ergodicity in the MBL phase—primarily via a *runaway avalanche instability* or, not necessarily independently, through long-range resonances.

Thermal avalanches

Firstly introduced in Ref. [243], the authors proposed that the MBL phase can become unstable in the presence of anomalously weakly disordered regions within the system. These regions may act as 'thermal bubbles', that expand through the neighboring degrees of freedom rendering them ergodic.

Consider a disordered spin chain of total length L, that contains a subsystem B (of length L_B) where the disorder is small and can be locally described by the disorder width W_B . The rest of the chain is described by the disorder W_A and both chains interact through the interface between A and B. The total Hamiltonian of the system $\hat{\mathcal{H}}$ can be expressed as

$$\hat{\mathcal{H}} = \hat{\mathcal{H}}_A + \hat{\mathcal{H}}_B + \hat{\mathcal{H}}_{AB} , \qquad (3.101)$$

in which $\hat{\mathcal{H}}_A$ is in the MBL phase (admitting a LIOM description), $\hat{\mathcal{H}}_B$ is the ergodic or thermal 'bubble' (that obeys ETH) and $\hat{\mathcal{H}}_{AB}$ involves local operators with support in both MBL and ETH subsystems, this is represented in a diagram in Fig. 3.7.

We prepare the system in the state

$$|\psi\rangle = |I_1 \dots I_j \dots I_{L_A}\rangle \otimes |n_B\rangle$$
, (3.102)

where $|I_1 \dots I_j \dots I_{L_A}\rangle$ denotes a configuration of the LIOMs³, i.e., a simultaneous eigenstate of the operators \hat{I}_k , with \hat{I}_k the LIOM centered at site k. In this formulation, the LIOMs correspond to effective spin- $\frac{1}{2}$ degrees of freedom, so that $I_j \in \{\uparrow, \downarrow\}$ specifies the local eigenvalue at site j. The state $|n_B\rangle$ is a many-body eigenstate of the thermal bubble Hamiltonian $\hat{\mathcal{H}}_B$.

³Here the LIOMs are defined in real space, rather than in terms of the Anderson orbitals as in Sec. 3.2.5. The two descriptions are equivalent up to a basis change.

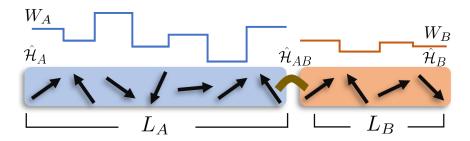


Fig. 3.7: The avalanche scenario. Subsystem A is subject to strong disorder and it is in the MBL regime. The subsystem B is weakly disordered and it is in the ergodic regime. The two subsystems interact via the boundary term \mathcal{H}_{AB} , which may induce a gradual thermalization of the subsystem B. Figure adapted from Ref. [254].

In first-order perturbation theory, the rate for the LIOM at site k to flip out of its initial eigenstate due to the coupling with the thermal bubble is given by the Fermi golden rule:

$$\Gamma_k \propto \rho_B e^{-2r/\xi_{\rm op}}$$
, (3.103)

where ρ_B is the density of states of the bubble, r is the distance between site k and the closest site belonging to the thermal bubble, and $\xi_{\rm op}$ is the localization length of the LIOM operators. The corresponding thermalization time $\tau_k \sim 1/\Gamma_k$ therefore increases exponentially with r.

To thermalize the LIOM closest to the bath i.e. $I_{L_A} \to \tilde{I}_{L_A}$ —where \tilde{I}_{L_A} indicates the reverse configuration e.g. if it was associate to a spin degree of freedom \uparrow now it becomes \downarrow . Then the thermalization time is $\tau_{L_A} \sim \rho_B e^{2/\xi_{\rm op}}$.

According to the avalanche mechanism [243], once the state of the LIOM \hat{I}_{L_A} has relaxed, the spin at site $j=L_A$ becomes a member of the ergodic bath, and, consequently the density of states the new bath B' is increased by a factor of 2, $\rho_{B'}=2\rho_B$. After r iterations of this process, the density of states of the enlarged bath B' is $\rho_{B'}=2^r\rho_B$.

As a consequence, thermalization of the MBL subsystem can become a self-sustaining process—i.e., the avalanche mechanism sets in—if the growth of the bath's density of states outpaces the exponential decay of the LIOM–bath coupling with distance r. To first order in perturbation theory, this condition reads

$$\frac{\ln 2}{2} - \frac{1}{\xi_{\rm op}} > 0 , \qquad (3.104)$$

which means that the rate of density-of-states growth of the expanding bath exceeds the rate at which the coupling strength decays with distance.

In summary, the avalanche scenario predicts that when the localization length of the LIOMs, $\xi_{\rm op}$, exceeds the critical value $\xi_{\rm op}^*=1/\ln 2$, the growth of the bath's density of states outpaces the exponential decay of the LIOM–bath coupling, enabling rare ergodic inclusions to destabilize the MBL phase and drive a thermalizing front through the system. For $\xi_{\rm op} < \xi_{\rm op}^*$, such avalanches cannot propagate and localization is expected to persist. In one dimension, the balance between coupling decay and density-of-states growth can be satisfied, whereas in higher dimensions avalanches are expected to proliferate generically, precluding a stable MBL phase.

In the thermodynamic limit, the scenario relies on the existence of *rare thermal bubbles*—small regions of anomalously weak disorder that act as ergodic grains capable of seeding the avalanche. In finite-size numerical simulations, however, the probability of finding such bubbles spontaneously is exponentially small, making the mechanism effectively invisible in typical disorder realizations. To overcome this, *planted-bubble* protocols are used: a thermal

grain is intentionally introduced by setting a contiguous region to have low or no disorder, embedded in an otherwise localized system.

This planted bubble mimics the role of a rare thermal inclusion and allows one to study how its thermalizing influence penetrates the MBL region. The growth of the *thermalization length* $\xi_d(t)$ can then be monitored, providing a practical test of avalanche theory in numerically accessible sizes [252, 254, 258]. Results from such studies are consistent with avalanche predictions, such as logarithmic growth of $\xi_d(t)$, but cannot fully rule out alternative microscopic mechanisms that produce similar dynamics. Thus, the problem of pin-pointing the microscopic mechanism responsible for the emergence of ergodicity in disordered spin chains remains a challenge.

Many-body resonances

Many-body resonances occur when an eigenstate of a many-body system can be expressed as a superposition of a few not entangled or weakly entangled states, e.g., the basis states $|I\rangle_S$. As it was explained in the beginning of Sec. 3.2, this implies that the basis states are nearly-degenerate and transitions can occurs between them, making it possible for the system to go flip a number of spins in the chain. If the resonance is system-wide (or long-range) meaning, that connects two basis states with a large amount of differing number of spins, then the system can spread outside of the range expected from the local integrals of motion. The proliferation of many-body resonances leads to the destabilization of MBL and the onset of ergodicity [264–268].

In Ref. [265] Villalonga and Clark used numerically constructed local integrals of motion to investigate many-body resonances and found that, in the vicinity of the ETH–MBL crossover, resonances occur across all length scales. This highlights that the crossover regime is characterized by widespread hybridization of eigenstates, in contrast to the fully ergodic or fully localized phases. Building on this perspective, Morningstar *et al.* [257] introduced the measure of quantum mutual information (QMI) between the first and last spins of a disordered XXZ spin chain (with open boundary conditions) as a diagnostic of such long-range resonances.

The QMI is particularly insightful because it captures correlations between the two ends of the system: it becomes small in both limiting regimes—ergodic and localized. In the ergodic phase, entanglement entropy obeys volume-law scaling, so the mutual information between two single spins essentially cancels out and vanishes. In the many-body localized (MBL) phase, end spins are uncorrelated due to localization, so the QMI again vanishes. However, in the intermediate crossover regime, the QMI becomes enhanced due to the superpositions of a few basis states with substantially different spin configurations—normally associated with 'cat-like states' ⁴. As such, QMI serves as a sensitive probe of system-wide resonances.

Morningstar *et al.* [257] quantified this effect by identifying, for each eigenstate, the maximum QMI across the spectrum and then taking the typical value of this quantity over disorder realizations. When plotted as a function of disorder strength W for different system sizes L, these curves display a crossing point at a characteristic scale $W_{\rm swr}(L)$. Interestingly, this crossing occurs at significantly larger disorder strengths than one predicted from standard observables $W^*(L)$. For example, at system size L=16, one finds $W^*\simeq 3.5$ while $W_{\rm swr}\simeq 8.5$. Moreover, as the system size is increased (e.g., from L=11 to L=16, $W_{\rm swr}(L)$ drifts upward by about one-third of its value, comparable to the drift observed for $W^*(L)$.

⁴Strictly speaking, an ideal Schrödinger cat state refers to a coherent superposition of two macroscopically distinct classical configurations. In the MBL crossover, the situation is more nuanced: the hybridizing states are not always perfectly classical, may involve more than two configurations, and often differ only over part of the system. The term 'cat-like' is therefore used heuristically to describe these long-range resonances rather than an exact two-state superposition.

Although the resonances detected in this way are extremely weak—the QMI at the crossing point is on the order of 10^{-7} —they nonetheless span the entire system, reflecting correlations between the two ends of the chain. This indicates a subtle form of non-locality emerging in the crossover regime. Furthermore, the maximum of the QMI as a function of W provides an additional marker for the MBL crossover, occurring at a disorder strength comparable to $W^*(L)$. Other indicators where studied in this study, such as the minimal value of the gap ratio $\langle r_{\rm min} \rangle$. The latter, when normalized to Poisson statistics, defines another disorder width $W_{\rm mg}(L)$ lying between the scales identified by W^* and $W_{\rm swr}$ around $\simeq 5.7$.

These observations motivate a careful study of the role of system-wide resonances in the MBL transition. In particular, extreme value statistics appear to be a natural tool to probe their impact, improving upon the standard observables discussed above. Understanding the effect of rare, long-range resonances is crucial for assessing the stability of the MBL phase and the nature of the transition, and the present study is aimed in this direction.

3.5 The large deviation method

3.5.1 The delocalization probability

A suitable order parameter for detecting delocalization is the probability that a system, initially prepared in the basis state $|0\rangle$, is found in distant basis states $|f\rangle$ after infinite-time.

We select the initial basis state $|0\rangle$ to be in the middle of the many-body spectrum, so as to probe localization in this highly-excited regime. This is done by selecting the diagonal elements of the Hamiltonian, that correspond to the random part $\hat{\mathcal{H}}_0$ in the spin basis. These energies $E_0 = \langle 0|\hat{\mathcal{H}}_0|0\rangle$ are extensive and normal distributed with zero mean and a variance proportional to L. As a result, $E_0/L \sim 1/\sqrt{L}$.

Basically, we want to sudy the long-time spreading of the wave-packet starting from $|0\rangle$, and examine whether it can reach distant configurations on the Hilbert space graph. For most of the work reported here, and unless otherwise stated, we will consider the distant basis states $|f\rangle$ to be states maximally uncorrelated with the initial one $|0\rangle$. The criteria of selection for these 'distant' or maximally uncorrelated configurations will be based on the overlap between the basis states $|0\rangle$ and $|f\rangle$, which depends on the choice of basis. In the spin basis the overlap is defined as

$$q_{0f}^{S} = \frac{4}{L} \sum_{i=1}^{L} S_{i}^{z}(0) S_{i}^{z}(f) , \qquad (3.105)$$

where $S_i^z \in \{\pm \frac{1}{2}\}$ is the eigenvalue of the spin operator in the z direction \hat{S}_i^z acting on the i-th site of the chain. Whereas, in the Anderson basis the basis states are occupation numbers, i.e. bit-strings made of 1s and 0s and the overlap reads

$$q_{0f}^{A} = \frac{1}{L} \sum_{\alpha=1}^{L} (2n_{\alpha}(0) - 1)(2n_{\alpha}(f) - 1) , \qquad (3.106)$$

where n_{α} is the occupation number of the α -th single-particle orbital, i. e. the eigenvalues of the $\hat{b}_{\alpha}^{\dagger}\hat{b}_{\gamma}\hat{b}_{\alpha}$ operators.

In both cases, $-1 \le q_{0f}^{S,A} \le 1$, and basis states that are completely uncorrelated with the initial state are characterized by zero overlap. In the spin basis this corresponds to flipping half of the spins between 0 and f. In the Anderson basis this implies that half of the L/2 spinless fermions have transitioned to different single-particle orbitals. Such overlaps are

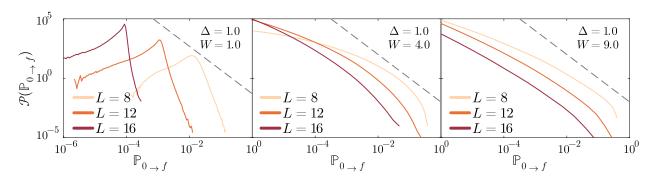


Fig. 3.8: Probability distribution function for the infinite time probabilities, Eq. (3.108), of finding a system, initially prepared in the basis state $|0\rangle$, in a basis state $|f\rangle \in \mathcal{E}$. The gray dashed line indicates a reference power-law decay with exponent 2. For weak disorder (W=1, left panel), the distribution is relatively narrow. As the disorder strength increases (W=4 and W=9, center and right panels), the distribution broadens significantly. The total number of samples used to compute these distributions is $N_{\text{tot}} = N_0 \times N_S \times \mathcal{N}_{\mathcal{E}}$, where $N_0 = 2^{L/2-2}$, and $N_S = 500, 5 \times 10^3, 5 \times 10^4$ for L=8,12,16, respectively.

closely related to the concept of *imbalance*, a measurement frequently used in the study of many-body quantum systems to quantify the degree to which memory of an initial basis state is retained over time [102, 406]. Based on the overlap value q, with respect to the initially prepared basis state $|0\rangle$, we define two sets of basis states:

$$S_0(q) = \{ |f\rangle : q_{0f}^S = q \} ,$$

$$A_0(q) = \{ |f\rangle : q_{0f}^A = q \} ,$$
(3.107)

for the spin and Anderson bases, respectively. In the following, we will primarily focus on $\mathcal{S}_0(q=0)$ and $\mathcal{A}_0(q=0)$, which we will collectively denote by the symbol \mathcal{E} and refer to as the 'equator' of the Hilbert space graph. It will be clear from the context whether \mathcal{E} refers to $\mathcal{S}_0(q=0)$ or $\mathcal{A}_0(q=0)$. Whenever ambiguity arises, we will explicitly add the superscript S or A to indicate that the quantity is measured in the spin or Anderson basis, respectively. In the zero-magnetization (or half-filling) sector, the equator possesses $\mathcal{N}_{\mathcal{E}} = \binom{L/2}{L/4}^2 \sim 2^L/(\pi L/4)$ target basis states, for any initial condition $|0\rangle$ chosen. The fact that the volume of the equator of the graph scales asymptotically with L in the same way as the volume of the entire graph, $\mathcal{N}_{\mathcal{E}}/\mathcal{N} \sim \sqrt{8/(\pi L)}$, is a hallmark of infinite-dimensional geometries. This justifies our primary choice of $q_{0f}^{S,A}$ for the analysis of localization: If the wave packet can reach \mathcal{E} in the infinite-time limit, it can essentially reach any configuration on the Hilbert space graph that is compatible with energy conservation.

The asymptotic probability that a system prepared in $|0\rangle$ reaches a basis state in the equator, $|f\rangle \in \mathcal{E}$, can be calculated exactly in terms of the eigenstates $|n\rangle$ of the Hamiltonian $\hat{\mathcal{H}}$ as

$$\mathbb{P}_{0\to f} = \lim_{t\to\infty} |\langle f|e^{-i\hat{\mathcal{H}}t}|0\rangle|^2 = \sum_{n} |\langle f|n\rangle\langle n|0\rangle|^2.$$
 (3.108)

Our key observable of interest is the overall *delocalization probability* that quantifies the likelihood that a system, initially prepared in the state $|0\rangle$, evolves into *any* state belonging to the equator:

$$\mathbb{P}_{\mathcal{E}} = \sum_{f \in \mathcal{E}} \mathbb{P}_{0 \to f} \ . \tag{3.109}$$

The sum runs over an exponentially large number of terms in L. This object is a random variables that depends on the disorder realization and on the choice of the initial state. In

a localized phase, the typical value of the delocalization probability should vanish with increasing system size, while go to 1 in the extended phase.

Using exact diagonalization of the full spectrum, we calculated the probability density function (PDF) of $\mathbb{P}_{0 \to f}$ for finite-sized systems for $\Delta = 1$. The results are shown in Fig. 3.8. In the inset we present a close-up of the right tails of the PDF of $\mathbb{P}_{0 \to f}$ in the spin basis. To construct this PDF, we selected N_0 initial states—located near the center of the energy spectrum—for each of the N_S disorder realizations of the random fields $\{h_i\}$. For each initial state, we computed $\mathbb{P}_{0 \to f}$ for all $\mathcal{N}_{\mathcal{E}}$ basis states with zero overlap from the initial one (i.e., at the equator of the Hilbert space graph). Consequently, each pdf is built from a total of $N_S \times N_0 \times \mathcal{N}_{\mathcal{E}}$ data points. The specific values of N_S , N_0 , and $\mathcal{N}_{\mathcal{E}}$ used for each system size are provided in the caption of Fig. 3.8.

For weak disorder (W=1), the distribution is sharply peaked. Consequently, the sum Eq. (3.109) is primarily governed by the bulk of the distribution $\mathbb{P}_{0\to f}$, with dominant contributions coming from its peak. This makes $\mathbb{P}_{\mathcal{E}}$ a self-averaging quantity. The peak of the distribution of $\mathbb{P}_{0\to f}$ shifts toward smaller values as the system size L increases. This indicates that as L is increased an exponentially increasingly large number of terms must contribute to the sum Eq. (3.109) in order to have that $\mathbb{P}_{\mathcal{E}}$ is of order 1.

As disorder is increased, instead, the tails of the distribution of $\mathbb{P}_{0\to f}$ decay increasingly slowly. At strong enough disorder this decay becomes slower than a square power-law, shown as a reference with a dashed-gray line. As a result, the sum in Eq. (3.109) becomes dominated by contributions from the tails of the distribution—that is, by rare events. These rare outliers from the tails correspond precisely to the system-wide strong resonances between states $|0\rangle$ and $|f\rangle$, which have been extensively discussed in the recent literature [257, 258, 266, 267, 270, 272, 277, 278].

The properties of $\mathcal{P}(\mathbb{P}_{0\to f})$ directly reflect on those of $\mathcal{P}(\mathbb{P}_{\mathcal{E}})$. In Fig. 3.9, we present these PDFs for four different disorder strengths: W=1, 4, 9, and 20, with system size L=16 and interaction strength $\Delta=1$. For W=1, $\mathbb{P}_{\mathcal{E}}$ is strongly peaked, with rapidly decaying tails. This is better appreciated in the inset, where we show a zoom-in with respect to the $\mathbb{P}_{\mathcal{E}}$ axis. As the disorder strength increases the PDFs become increasingly broad and

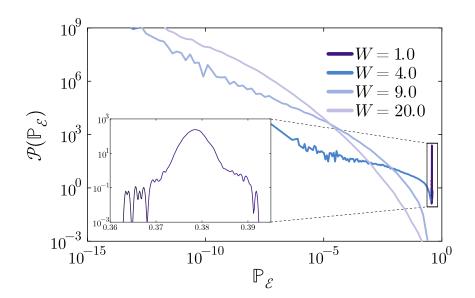


Fig. 3.9: Probability distributions of the delocalization probability $\mathbb{P}_{\mathcal{E}}$ in log-log scale, for the disorder strengths shown in the key, with L=16 and $\Delta=1$. The inset is a zoom-in for the W=1 distribution, shown in linear-log scale instead. This latter distribution is heavily peaked at finite values of $\mathbb{P}_{\mathcal{E}}$ with fast decaying tails.

asymmetric: most values cluster near a very small value; however, the distribution develops increasingly heavy tails, which are rare realizations of the disorder for which the probability of delocalizing is anomalously large. In the inset, one can observe that the right tails for W=4 and W=9 even extend beyond the full distribution for W=1. This indicates that some disorder realizations at very strong disorder strength have the same probability to delocalize to that of the typical samples at disorder W=1.

In this work, we argue that these rare events are responsible for the presence of ergodic instabilities inside the MBL at finite sizes, within a broad intermediate region of the phase diagram. The central goal of the method proposed here is to put forward a computational scheme that allows one to account correctly for the statistical weight of the fat-tailed distributions $\mathcal{P}(\mathbb{P}_{0\to f})$. Such a computational scheme is based on an analogy with classical mean-field disordered systems that undergo phase transitions exhibiting phenomenological features similar to those described above for the probability of delocalization from a random initial state. Specifically, the partition function of a classical disordered system with N degrees of freedom is given by

$$\mathcal{Z}_N = \sum_{\mu} e^{-\beta E_{\mu}} \,, \tag{3.110}$$

where the μ 's label microscopic configurations of the system, whose number grows exponentially with N. In general, the energies E_{μ} are random (and correlated). The inverse temperature β controls the spread of the Boltzmann weights. In this analogy, the partition function \mathcal{Z}_N corresponds to the delocalization probability $\mathbb{P}_{\mathcal{E}}g$, while the Boltzmann factors $e^{-\beta E_{\mu}}$ play the role of the transition probabilities $\mathbb{P}_{0\to f}$.

A broad class of classical mean-field disordered systems exhibit a sharp phase transition in the thermodynamic limit $(N \to \infty)$, from a high-temperature phase—where the partition function receives contributions from an exponential number of configurations—to a low-temperature phase, where the Boltzmann measure freezes onto a few rare configurations with anomalously large weights in the tails of the Boltzmann factor distribution [407–412].

The key result for such systems is that, in the frozen phase, in the thermodynamic limit the *typical* value of the partition function (which corresponds to the free-energy in classical disordered systems terminology) is dominated by rare configurations in the far tails of the Boltzmann weight distribution. In the MBL context, this implies that the *typical* value of the delocalization probability $\mathbb{P}_{\mathcal{E}}$ is asymptotically controlled by a few anomalously strong, system-wide resonances.

Yet, accurately estimating the asymptotic typical value of the partition function (or equivalently $\mathbb{P}_{\mathcal{E}}$) from finite-size numerical simulations is extremely challenging, since accessible system sizes typically do not include the rare events that dominate the measure in the thermodynamic limit, resulting in strong finite-size corrections.

Our approach therefore consists in adapting the set of tools and methods developed in the study of mean-field classical disordered systems to properly account for these rareevent effects, allowing us to evaluate the correct asymptotic behavior of the delocalization probability in a disorder regime where it is given by a sum over an exponentially large number of correlated and broadly distributed random variables.

However, the computational cost of evaluating Eq. (3.109) is substantial, specifically, it requires computing all eigenstates $|n\rangle$ of the Hamiltonian $\hat{\mathcal{H}}$. Carrying out exact diagonalization of the full spectrum over a sufficiently large number of disorder realizations to obtain reliable statistics is computationally intensive and can only be done for relatively small sizes, $L \leq 16$ (corresponding to a maximal Hilbert space dimension of $12\,870$). To address this limitation, we introduce a proxy quantity for $\mathbb{P}_{0\to f}$ that is easier to access computationally, yet still retains the same physical information. This allows us to capture the essential features of the transport properties under investigation. An additional advantage of this approach is

that it enables the exploration of larger system sizes, up to L=22, which corresponds to a Hilbert space dimension of $705\,432$.

3.5.2 A Proxy for Transport in Hilbert space

Let us now go back to the set-up defined in Sec. 3.5.1, in which we probe how a system, initially prepared in the random basis state $|0\rangle$, spreads to an exponentially large subset of basis states located at the equator \mathcal{E} . In this setting, Eqs. (3.108) and (3.109) give the definition of the probability to decorrelate from a random initial state after infinite time. In a many-body system, the projections of the eigenstates $|n\rangle$ with energy densities different from the initial energy $\langle 0|\mathcal{H}|0\rangle$ are expected to vanish exponentially with the system size. As a result, the sum over $|n\rangle$ is expected to be dominated by the eigenstates having the same expectation value of energy as the initial state. Hence, the probability that the system is initialized in $|0\rangle$ and is found in $|f\rangle$ after an infinite time can be expressed in terms of the squared modulus of the infinite-time propagator between the nodes of the Hilbert space corresponding to the basis states $|0\rangle$ and $|f\rangle$:

$$\mathbb{P}_{\mathcal{E}} \approx \lim_{\eta \to 0} \sum_{f \in \mathcal{E}} \eta |\mathcal{G}_{0f}(E - i\eta)|^2 , \qquad (3.111)$$

where \mathcal{G}_{0f} are the off-diagonal elements of the resolvent $\hat{\mathcal{G}}(E-\mathrm{i}\eta)=(E-\mathrm{i}\eta-\hat{\mathcal{H}})^{-1}$ computed on the nodes $|0\rangle$ and $|f\rangle$ of the Hilbert space, and η is a small imaginary regulator.

The interpretation of Eq. (3.111) is quite intuitive: The spreading of the many-body states is driven by energy-resonant hybridization among basis states with energies close to the selected energy E. In closed systems, such resonances are captured by the off-diagonal elements of the resolvent operator, whose matrix elements, $\mathcal{G}_{0f} = \langle f | \hat{\mathcal{G}} | 0 \rangle$, quantify the effective hopping amplitude for an energy-resolved transition between the initial state $|0\rangle$ and the target state $|f\rangle$.

As detailed in Appendix A1, computing the matrix elements \mathcal{G}_{0f} is computationally more efficient than evaluating the probabilities $\mathbb{P}_{0\to f}$, as it only requires solving a sparse linear system (see Eq. (141)), rather than obtaining the full energy spectrum. Nevertheless, \mathcal{G}_{0f} encodes the same physical information as $\mathbb{P}_{0\to f}$.

Specifically, Eq. (3.111) can be interpreted in terms of an analogy from the study of quantum mesoscopic systems in real space, where transport properties at fixed energy E are typically characterized by the Landauer transmission [413]. These are commonly computed using the Fisher-Lee formula [289, 414], which relates the dimensionless conductance to the Green's function of the scatterer dressed by the incoming and outgoing leads—creating channels of transport. Explicitly, the formula reads

$$\mathcal{T}_{FL} = \text{Tr}\{\Gamma_L \mathcal{G}^r \Gamma_R \mathcal{G}^a\} , \qquad (3.112)$$

where the energy dependence E is implicit in all quantities. The superscripts r and a denote the retarded and advanced Green's functions, respectively, which are related through $\mathcal{G}^a = (\mathcal{G}^r)^{\dagger}$. The quantities Γ_L and Γ_R represent the level broadening due to the coupling with the left (L) and right (R) leads, respectively, and are given by

$$\Gamma_{L,R} = -2 \operatorname{Im} \Sigma_{L,R} , \qquad (3.113)$$

with $\Sigma_{L,R}$ being the self-energies of the leads. A schematic representation of this construction is shown in Fig. 3.10(a). In this respect, $\mathbb{P}_{\mathcal{E}}$ is the analogue of the conductance of a complex network (i.e., the Hilbert space graph) in a scattering geometry in which a semi-infinite lead through which 'particles' are injected is connected to the node $|0\rangle$, and $\binom{L/2}{L/4}^2$ semi-infinite

leads through which 'particles' are extracted are connected to the nodes $|f\rangle \in \mathcal{S}_0(0)$ (or $\mathcal{A}_0(0)$, depending on the chosen basis). This construction is illustrated in Fig. 3.10(b), where the incoming and outgoing semi-infinite leads are replaced by the initial state $|0\rangle$ and the set of target states \mathcal{E} , respectively. It is clear that in this analogy the broadening of the energy levels provided by the imaginary part of the self-energy of the leads plays the role of the small imaginary regulator η in Eq. (3.111). In the ergodic phase—where all channels contribute comparably to transport— $\mathcal{T}_{\mathrm{FL}}$ is of order 1. By contrast, in the MBL regime transport is suppressed and the typical value of $\mathcal{T}_{\mathrm{FL}}$ decays exponentially with L.

In order to further simplify the numerical computations, in most of the following we do not evaluate the true delocalization probability Eq. (3.109) (or, equivalently, the Landauer transmission Eq. (3.112)), but rather a simplified proxy in which the imaginary parts are neglected:

$$\mathcal{T}_0 = \sum_{f \in \mathcal{E}} |\mathcal{G}_{0f}|^2 , \qquad (3.114)$$

where $\mathcal{G}_{0f} = \langle f | (E - \hat{\mathcal{H}})^{-1} | 0 \rangle$ denotes the real-part off-diagonal propagator between $| 0 \rangle$ and $| f \rangle$. Restricting the computation to real parts only reduces the numerical cost—both in terms of computational time and, more importantly, memory usage—by more than a factor of two.

Yet, in the absence of the imaginary regulator, the amplitudes $|\mathcal{G}_{0f}|^2$ are no longer bounded and can take arbitrarily large values. This occurs because the poles of \mathcal{G} associated with many-body eigenstates whose energies lie very close to E are no longer regularized, leading to spuriously large contributions. As a result, the $|\mathcal{G}_{0f}|^2$'s, and hence \mathcal{T}_0 , can no longer be strictly interpreted as the probabilities to delocalize from $|0\rangle$ to $|f\rangle$ after infinite time.

The absence of this regularization has, in particular, important consequences for the asymptotic scaling of \mathcal{T}_0 in the delocalized phase. While the properly regularized transmission $\mathcal{T}_{\mathrm{FL}}$ and the probability to reach the equator $\mathbb{P}_{\mathcal{E}}$ saturate to a value of order one on the ergodic side of the transition, their unregularized counterpart \mathcal{T}_0 grows exponentially with the system size L, proportionally to the number of outgoing channels. In contrast, in the MBL phase both $\mathcal{T}_{\mathrm{FL}}$ and \mathcal{T}_0 exhibit the same exponential decay with L. Therefore, although \mathcal{T}_0 formally overestimates the true Landauer transmission, it reproduces its correct large-L scaling in the localized phase and can thus be reliably (and more efficiently) used to identify the localization transition, as discussed below.

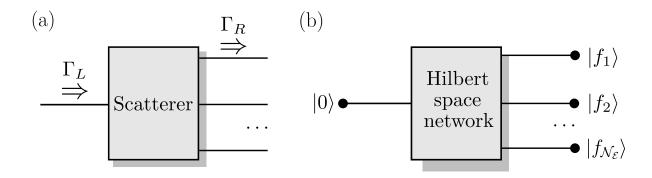


Fig. 3.10: (a) Quantum transport on a network in a scattering geometry, receiving particles from a semi-infinite lead on the left and transmits them through several semi-infinite leads connected to its right-hand side. (b) Schematic of the transport of the 'fictitious particle'—initially prepared in the basis state $|0\rangle$ —on the Hilbert space network.

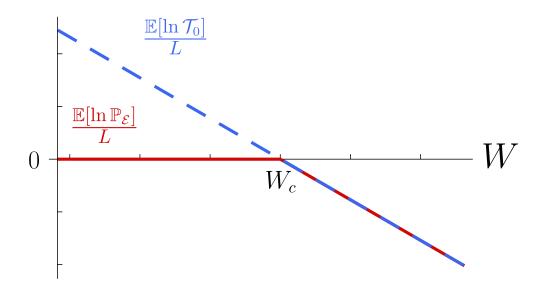


Fig. 3.11: Sketch of the different scaling behavior with L of the typical value of the Landauer transmissions with and without imaginary parts for the Anderson model on the Bethe lattice.

Here, we present several arguments to support this claim and justify our choice. The first argument stems from an analogy with the Anderson model on the Bethe lattice. Indeed, the Hilbert-space graph associated with a quantum many-body problem—which lies at the core of our discussion—is, in general, a sparse and high-dimensional network, as discussed in Sec. 3.3.1. Although this graph exhibits strong correlations and complex loop structures that are absent in simpler tree-like models such as the Bethe lattice, the latter nonetheless serves as a valuable toy model that captures the essential features of the Hilbert-space network [349]. This analogy provides a powerful framework for gaining qualitative insight into the behavior of many-body systems [283, 289, 350–352, 415–417].

The order parameter for the localization transition in the single-particle Anderson model is the typical value of the imaginary part of the local Green's function, $\exp\{\mathbb{E}[\ln \operatorname{Im} \mathcal{G}_{ii}]\}$, which is finite in the delocalized phase and vanishes in the localized one. For locally tree-like graphs, asymptotically exact recursive relations can be derived that express these local Green's functions in terms of those on neighboring nodes [418]. Considering a specific node 0 of the tree, the recursive equation for the imaginary part (in the $\eta \to 0$ limit) can be telescoped as:

$$\operatorname{Im} \mathcal{G}_{00} = \left| \mathcal{G}_{00} \right|^2 \sum_{i \in \partial 0} \operatorname{Im} \mathcal{G}_{ii}, \qquad (3.115)$$

where the sum runs over all neighbors i of node 0, and $|\mathcal{G}_{00}|^2$ is the squared modulus of the full local Green's function, including both its real and imaginary parts. Starting from node 0, we iteratively apply the recurrence to unfold the expression, rewriting the imaginary parts $\operatorname{Im} \mathcal{G}_{ii}$ on the right-hand side in terms of those from successive generations of the tree, i.e., of nodes at increasing distance from 0. Repeating this procedure over L generations, and writing the two-points propagators as product of the local Green's functions yields (see Refs. [281, 419] for more details):

$$\operatorname{Im} \mathcal{G}_{00} = \sum_{f=1}^{k^L} |\mathcal{G}_{0f}|^2 \operatorname{Im} \mathcal{G}_{ff}, \qquad (3.116)$$

where f labels the k^L nodes at distance L from 0 (where k is the branching ratio of the tree). This equation is the analogue of Eqs. (3.111) and (3.112), and its physical interpretation is

intuitively clear: Im \mathcal{G}_{00} represents the inverse lifetime of a particle created at node 0, while the propagators $|\mathcal{G}_{0f}|^2$ correspond to the probabilities that the particle escapes from 0 to f after an infinite time.

We now set the imaginary part of the Green's functions on all these distant nodes f to a small value η , and take L to be large. In the delocalized phase, the imaginary parts grow under iteration and eventually reach a stationary value for large L. In the localized phase, by contrast, they decay exponentially under iteration. In this regime, the recursion relations can be linearized with respect to the imaginary parts, and to leading order Eq. (3.116) becomes:

Im
$$\mathcal{G}_{00} = \eta \sum_{f=1}^{k^L} |\text{Re } \mathcal{G}_{0f}|^2 = \eta \, \mathcal{T}_0 \,.$$
 (3.117)

The exponential decay of the typical value of $\operatorname{Im} \mathcal{G}_{00}$ under iteration in this linearized regime is governed by the largest eigenvalue of an integral operator that encodes the critical properties of Anderson localization [355, 418]. One finds $\exp\{\mathbb{E}[\ln\operatorname{Im} \mathcal{G}_{00}]\} \simeq \eta \, \lambda_{\max}^L$, with $\lambda_{\max} \simeq 1 - c(W - W_c)$ near the localization transition, hence

$$\frac{1}{L}\mathbb{E}[\ln \mathcal{T}_0] = \ln \lambda_{\max} \simeq -c(W - W_c). \tag{3.118}$$

Consequently, the scaling of the typical Landauer transmission with L (evaluated from the real propagators) reflects the W-dependence of the Lyapunov exponent that controls the response of the typical $\operatorname{Im} \mathcal{G}$ to perturbations, thereby indicating whether the system is in a localized or a delocalized phase. Note that in the context of the Anderson model on tree-like graphs, this is an exact result.

Of course, in the delocalized phase the recursion equations can no longer be linearized, and Eq. (3.117) no longer describes the probability of delocalization from a random initial state: instead of saturating to a finite value, as its counterpart (3.116) which contains also the imaginary parts, it diverges exponentially. Nevertheless, as discussed below and explicitly demonstrated in [281], the scaling of its typical value with L can still be used effectively and reliably to locate the localization transition in the Anderson model on the Bethe lattice. This argument is schematically depicted in Fig. 3.11 which illustrates the different behavior of the scaling with L of the typical value of the Landauer transmissions—or equivalently, the delocalization probability—with and without imaginary parts for the Anderson model on the Bethe lattice.

The second argument supporting the choice of the typical value of \mathcal{T}_0 as an order parameter for MBL comes from the benchmark analysis presented in Ref. [281]. In that work, some of us employed this quantity not only to locate the localization transition in the Anderson model on the Bethe lattice (for which, as discussed above, Eq. (3.118) can be explicitly and rigorously justified), but also to study abstract random matrix ensembles that display three distinct regimes: a fully delocalized phase, a localized phase, and an intermediate delocalized yet fractal phase. Ref. [281] demonstrated that the numerical scaling analysis of the typical value of \mathcal{T}_0 successfully reproduces the analytically known phase diagram of these models, accurately identifying the transitions between the different phases.

A final argument in favor of our approximation is provided by direct numerical tests [420] on the random-field transverse-field Ising model. We explicitly examined the scenario schematically illustrated in Fig. 3.11. In particular, we computed the probability to delocalize from a random initial state, using Eqs. (3.108) and (3.109) (via exact diagonalization), and compared it to the Landauer transmission \mathcal{T}_0 for the same disorder realizations and initial states. At sufficiently strong disorder—where $\mathbb{P}_{\mathcal{E}}$ decays exponentially with the system size—we find that the typical value of \mathcal{T}_0 is proportional to that of $\mathbb{P}_{\mathcal{E}}$. Moreover, their

covariance increases with system size and approaches unity in the strong-disorder regime for the system sizes accessible to our numerics [420]. This implies that samples and initial configurations with an anomalously large probability to decorrelate after infinite time also exhibit anomalously large values of \mathcal{T}_0 .

To summarize, even though neglecting the imaginary regulator in Eqs. (3.111) and (3.112) leads to a distinct scaling behavior in the delocalized phase (due to the proliferation of spurious poles in the denominator of \mathcal{G}), the typical value of the Landauer transmission \mathcal{T}_0 nonetheless provides an efficient and reliable order parameter for MBL. It offers a much simpler computational route while faithfully reflecting the same physical information: instead of directly evaluating the probability for a random initial state to delocalize, we compute a proxy quantity that signals whether this probability decays exponentially or not.

Yet, we are still left with the problem of how to perform a correct statistical analysis of \mathcal{T}_0 , which is a sum of an exponential number of correlated random variables, whose distribution becomes broader as W is increased.

As explained above, Eq. (3.114) is a sum of an exponential number of correlated random variables, whose distribution becomes broader and broader when W is increased. \mathcal{T}_0 is formally equivalent to the partition function of a classical disordered system (3.110).

Specifically, the factors $|\mathcal{G}_{0f}|^2$ play the role of the Boltzmann weights associated to each one of the exponentially numerous target states $|f\rangle$. As mentioned above, for a broad class of mean-field classical disordered systems, such as directed polymers in disordered media on high-dimensional graphs [421–423] and related models [407, 408, 411, 424], it is well known that the partition function (3.110) can undergo a phase transition: if the probability distribution of the Boltzmann weights becomes too broad (which in classical systems is induced by reducing the temperature), then the measure is dominated by few O(1) outliers of the distribution, corresponding to a few configurations of the system with particularly low energy; instead at high temperature the partition function receives significant contribution from exponentially many terms.

In this context, our ultimate goal is to estimate the typical value of \mathcal{T}_0 in the large-L limit, which is analogous to estimating the free-energy density in classical disordered systems at a given temperature. However, in the many-body problem, the broadening of the distribution of the $|\mathcal{G}_{0f}|^2$'s is indirectly controlled by the disorder strength W. For a given W, it is not obvious a priori whether \mathcal{T}_0 corresponds to the 'high-temperature' or 'low-temperature' phase in this analogy. This distinction is crucial, as estimating the asymptotic behavior of the typical value in the frozen phase is particularly subtle. As mentioned earlier, one must properly account for the contribution of rare outliers—namely, atypical disorder realizations that lead to anomalously large values of \mathcal{T}_0 and are unlikely to appear in small or moderate system sizes.

In the following, we begin by explaining how this type of analysis is carried out in the standard setting of classical disordered systems, where the distribution of Boltzmann weights is known and analytical calculations are possible. As a guiding example, we will focus on the directed polymer in disordered media [425], which has already been connected to various tightly-related quantum systems—most notably to single-particle Anderson localization [289, 426, 427], two-interacting-particle models in one-dimensional disordered systems [428], and random-field quantum Ising models [429, 430]. Following this discussion, we will outline how we adapt this theoretical framework to compute the typical value of \mathcal{T}_0 in our quantum many-body context.

The transition indicators

Freezing transition in mean-field classical disordered systems

The discrete model of a directed polymer in random media (DPRM) consists of a self-avoiding directed random walk on a d-dimensional lattice. Each edge of the lattice, (ij), has an associated energy ϵ_{ij} which is a quenched random variable. The set of edges that the self-avoiding random walk follows defines a path \mathcal{P} , and its energy is given by $E_{\mathcal{P}} = \sum_{(ij) \in \mathcal{P}} \epsilon_{ij}$. The partition function is

$$\mathcal{Z}_N(\beta) = \sum_{\mathcal{P} \in \mathcal{P}_N} e^{-\beta E_{\mathcal{P}}} , \qquad (3.119)$$

where $E_{\mathcal{P}}$ is the total random energy collected along the path \mathcal{P} of length N. The sum runs over all possible directed paths of length N, here denoted as the set \mathcal{P}_N , and β is the inverse temperature. The energies of different paths, say $E_{\mathcal{P}}$ and $E_{\mathcal{P}'}$, are correlated through their common edges. A formal relationship between this model and Anderson localization has been exploited in the past [289, 322, 427, 431, 432]. DPRM exhibit a well-known freezing transition in the infinite-dimensional limit [421, 423], when the problem is studied on an infinite tree. On this class of hierarchical lattices, the number of directed paths for polymers of length N is $(k+1)k^{N-1}$, k being the branching ratio of the tree.

The transition occurs at a critical inverse temperature $\beta = \beta_{\star}$. At high temperatures, $\beta < \beta_{\star}$, the partition function receives contributions from an exponential number of directed paths. Instead, in the low temperature phase $\beta > \beta_{\star}$, the polymer freeze in a few O(1) specific disorder-dependent paths. This corresponds to a condensation of the Boltzmann measure on a few paths with particularly low energy, which directly reflects in the non-analytic behavior of the *quenched* free-energy density, defined as

$$f_q(\beta) = -\lim_{N \to \infty} \frac{1}{\beta N} \langle \ln \mathcal{Z}_N(\beta) \rangle ,$$
 (3.120)

where $\langle \cdots \rangle$, denotes the average over the disorder realizations. For the specific problem of directed polymers on a tree, the exact solution of Refs. [421, 423] yields:

$$f_q(\beta) = \begin{cases} -\ln\left(k\left\langle e^{-\beta\epsilon}\right\rangle\right)/\beta & \text{for } \beta < \beta_{\star}, \\ -\ln\left(k\left\langle e^{-\beta_{\star}\epsilon}\right\rangle\right)/\beta_{\star} & \text{for } \beta \ge \beta_{\star}. \end{cases}$$
(3.121)

In the thermodynamic limit, the quenched free-energy density develops a plateau for $\beta > \beta_{\star}$. For finite system sizes, however, the quenched free-energy remains a concave function of β for any arbitrarily large but finite N. The convergence of the finite-N quenched free-energy to its asymptotic plateau value for $\beta > \beta_{\star}$ is slow, with finite-size corrections scaling as $\log N/N$ [433, 434]. This slow convergence arises because typical finite-size samples do not contain the rare configurations that dominate the behavior in the thermodynamic limit.

In this situation, it is particularly insightful to study also the behavior of the finite-size annealed free-energy density, defined as

$$f_a(\beta, N) = -\frac{1}{\beta N} \ln \langle \mathcal{Z}_N(\beta) \rangle$$
, (3.122)

and compare it with the finite-size quenched free-energy density (3.120).

The basic idea is the following: In the high-temperature phase ($\beta < \beta_{\star}$), the sum (3.119) is dominated by the bulk of the probability distribution of the Boltzmann weights. As a result, the average and typical values of \mathcal{Z}_N exhibit the same asymptotic scaling with N, and the free-energy is self-averaging, i.e., the annealed and quenched free-energy densities converge

to the same value upon increasing the system size N. In contrast, in the low-temperature phase ($\beta > \beta_{\star}$), only a few disorder-dependent paths with particularly low energies $E_{\mathcal{P}}$ dominate the sum in Eq. (3.119). These rare, large Boltzmann weights induce strong sample-to-sample fluctuations, leading to a broad distribution $P(\mathcal{Z}_N)$ characterized by power-law tails [421]. When these tails are sufficiently heavy, the typical and average values of \mathcal{Z}_N exhibit different scaling with N: the former is governed by the bulk of the distribution, while the latter is dominated by rare configurations with anomalously low energies. Consequently, the self-averaging property of the free-energy is lost. In particular, the extreme outliers skew the average $\langle \mathcal{Z}_N \rangle$, causing the finite-N annealed free-energy density to develop a maximum close to β_{\star} . Such maximum of the annealed free-energy curve is completely unphysical and purely a consequence of the biased finite sampling of anonymously large outliers of \mathcal{Z}_N .

Yet, in models for which the analytical solution in the thermodynamic limit is unavailable (as is the case for the asymptotic typical value of \mathcal{T}_0 in the MBL problem), the behavior of the annealed free-energy provides a valuable practical tool to estimate the asymptotic critical behavior of the quenched free-energy in the low-temperature phase using finite-N numerical results.

The basic idea is to first identify the position of the maximum of the annealed free-energy curve, which gives an estimate of β_{\star} . Then, the unphysical portion of the annealed free-energy for $\beta > \beta_{\star}$ is replaced by a flat segment at height $f_a(\beta_{\star}, N)$. For models in which the large-N solution is known analytically, this construction has been shown to provide a more accurate approximation of the asymptotic value of the quenched free-energy density in the frozen phase [281, 434]

Formally, this construction is as follows:

$$\tilde{f}_a(\beta, N) = \begin{cases}
f_a(\beta, N) & \text{for } \beta < \beta_{\star}, \\
f_a(\beta_{\star}, N) & \text{for } \beta \ge \beta_{\star},
\end{cases}$$
(3.123)

where $f_a(\beta, N)$ is the finite-size annealed free-energy density given in Eq. (3.122). With this modification one obtains two equivalent quantities for the entire temperature range in the thermodynamic limit:

$$f(\beta) = \lim_{N \to \infty} f_q(\beta, N) = \lim_{N \to \infty} \tilde{f}_a(\beta, N) . \tag{3.124}$$

At finite N, $f_q(\beta, N)$ and $\tilde{f}_a(\beta, N)$ are complementary, and allow us to assess the role that finite size effects play in the problem at hand. The former describes the behavior of typical samples at the chosen value of N; while the latter provides a more accurate estimation at large N [281].

Probing the rare events in the quantum many-body problem

In the following, we apply the construction (3.123) to estimate the asymptotic behavior of the typical value of \mathcal{T}_0 in the large-N limit, while properly accounting for the statistical contribution of strong, system-wide resonances which, for the system sizes accessible numerically, typically form only between a few specific distant configurations, and only in rare disorder realizations.

The key assumption underlying our approach is that Eq. (3.114) belongs to the same universality class as the classical mean-field disordered models discussed above. This connection can be rigorously established in the case of single-particle Anderson localization on the Bethe lattice [281]. More generally, this universality holds whenever the elements of the sum exhibit ultrametric correlations, as in the DPRM case [408, 410]. In Appendix A3,

we provide numerical evidence that the correlations $\langle |\mathcal{G}_{0f}||\mathcal{G}_{0f'}|\rangle_c$ are consistent with this assumption for the system sizes we are able to study.

To apply the machinery described in the previous section to the MBL problem, we enlarge the parameter space by formally introducing an auxiliary parameter β , which plays the role of the inverse temperature in the classical problem. We thus define the β -dressed version of the transmission, $\mathcal{T}_0(\beta)$:

$$\mathcal{T}_0(\beta) = \sum_{f \in \mathcal{E}} |\mathcal{G}_{0f}|^{\beta} . \tag{3.125}$$

This auxiliary parameter has the role of tuning the strength of the tails of the probability distribution of the $|\mathcal{G}_{0f}|$'s for a fixed value of the disorder W of the original MBL problem, thereby allowing us to identify the threshold value β_{\star} at which the distribution of $\mathcal{T}_{0}(\beta)$ develops sufficiently broad tails. At this point, rare outliers begin to dominate its typical value, signaling the onset of a freezing transition analogous to that in the DPRM. Note that $\mathcal{T}_{0}(\beta)$ has been already introduced and studied in the mathematical literature on single-particle Anderson localization in hierarchical lattices [435, 436], in the context of the so-called fractional moment method. Our primary interest is to compute the asymptotic scaling behavior of the typical value of $\mathcal{T}_{0}(\beta=2)$, which is a proxy for the probability tht a randomly chosen initial state reaches arbitrarily distant configurations on the Hilbert space graph after infinite time.

In order to do this, we have to determine the position of $\beta=2$ relative to the threshold β_{\star} , which determine whether $\mathcal{T}_0(\beta=2)$ is dominated by the broad tails of its distribution (i.e., rare events) or whether the tails decay rapidly enough for it to be dominated by the bulk of the distribution (i.e., typical realizations of the disorder).

Proceeding as in the theoretical framework outlined in the previous section, we define two quantities analogous to the annealed and quenched free-energy densities:

$$\phi_a(\beta, L) = \frac{\ln \mathbb{E}[\mathcal{T}_0(\beta, L)]}{\beta \ln \mathcal{N}_{\mathcal{E}}}, \qquad (3.126)$$

$$\tilde{\phi}_a(\beta, L) = \begin{cases}
\phi_a(\beta, L) & \text{for } \beta < \beta_{\star}, \\
\phi_a(\beta_{\star}, L) & \text{for } \beta \ge \beta_{\star},
\end{cases}$$
(3.127)

$$\phi_q(\beta, L) = \frac{\mathbb{E}[\ln \mathcal{T}_0(\beta, L)]}{\beta \ln \mathcal{N}_{\mathcal{E}}} , \qquad (3.128)$$

corresponding to the annealed, annealed with the plateau replacement — similar to Eq. (3.123) — and the quenched free-densities, respectively. Here $\mathbb{E}[\cdots]$ denotes the average with respect to initial conditions $|0\rangle$ and disorder realizations of the random fields.

The partition function of the classical DPRM $\mathcal{Z}_N(\beta)$, has been replaced by the biased Hilbert space Landauer transmission $\mathcal{T}_0(\beta)$. In the classical DPRM model, the total number of configurations—that is, the number of terms contributing to the partition function—scales exponentially with the length of the polymer as k^N , previously denoted by N. In this case, the number of terms contributing to the sum is $\mathcal{N}_{\mathcal{E}}$, which represents the number of target states. Accordingly, we adapt the definition of the free-energy functions by normalizing with $\ln \mathcal{N}_{\mathcal{E}}$. Additionally, we have omitted the overall minus sign in the definition, so that the free-energies become negative when the typical value of \mathcal{T}_0 vanishes—signaling localization—while a positive values of the free-energies indicates delocalization. Due to this modification the maximum in the annealed free-energy density of the original DPRM at β_* now appears as a minimum of $\phi_a(\beta)$. Beyond the dependence on β and λ , $\mathcal{T}_0(\beta, L)$ also depends on λ and λ . From this point onward, we will omit the explicit dependence on λ , λ and λ in the functions λ , λ , λ , and λ , and λ , λ , λ , and λ

Exploiting the analogy with DPRM and similar mean-field disordered models [408, 421, 423, 434, 437], the large-L behavior of the probability of delocorrelate from a random initial state, $\mathcal{T}_0(\beta=2)$, can be estimated indirectly in two steps, as described below:

- Ergodic Phase. We first determine the position of the minimum β_{\star} of the finite-L annealed free-energy density. If $\beta_{\star} > 2$ the Hilbert space Landauer transmission defined in Eq. (3.114) lies within the ergodic (or high-temperature) phase, and its typical value can be directly obtained from $\phi_q(\beta=2)$. In this regime the system is delocalized since $\phi_q(\beta=2) > 0$.
- Non-Ergodic Phases. Conversely, when $\beta_{\star} < 2$, physical transport occurs within the freezing region, where the Hilbert space Landauer transmissions are dominated by the tails of their distribution. According to our analogy with classical disordered systems and DPRM, in this regime, the finite-L value of $\phi_q(\beta=2)$ provides a poor estimate of the typical value of \mathcal{T}_0 in the large-L limit. This is because typical samples at small L lack the rare events that will dominate the statistics at large L. A more accurate estimation is instead provided by the value of $\tilde{\phi}_a(\beta=2) = \phi_a(\beta_{\star})$. Then:
 - Non-Ergodic Delocalized Phase. If $\phi_a(\beta_\star) > 0$, the elements in the tails of the distribution of $|\mathcal{G}_{0f}|^2$ give such a large contribution that the typical value of \mathcal{T}_0 remains positive upon increasing system size. This corresponds to a regime in which the system eventually delocalizes for large enough L through a small number of long-range resonances that destabilize the MBL phase.
 - Non-Ergodic Many-Body Localized Phase. If instead $\phi_a(\beta_\star) < 0$, even with the inclusion of large rare-events of \mathcal{T}_0 , the typical value of \mathcal{T}_0 vanishes with increasing system size. This corresponds to a genuine MBL phase, where rare system-wide resonances are unable to make a random initial state completely decorrelate.

These three distinct criteria provide us with independent estimates of finite-sized critical disorder strengths, separating the different regimes observed in the model:

- 1. $W_{\rm ergo}(L)$ is defined by the condition $\beta_{\star}(W_{\rm ergo})=2$. It provides an estimate for the disorder strength at which the system transitions from an ergodic regime—where physical transport at $\beta=2$ receives contributions from an exponential number of terms $|\mathcal{G}_{0f}|^2$ and its behavior is governed by typical instances—to a regime where \mathcal{T}_0 is dominated by few anomalously large outliers from the distribution's tails of $|\mathcal{G}_{0f}|^2$, and the distribution of \mathcal{T}_0 becomes broad.
- 2. $W_{\rm MBL}(L)$ is defined by the condition $\phi_a(\beta_\star, W_{\rm MBL}) = 0$. It provides an estimate of the critical disorder strength at which the system enters a genuinely localized regime. In this phase, even rare resonant inclusions between basis states are insufficient to induce delocalization.
- 3. $W_{\mathrm{MBL}}^{\mathrm{typ}}$ is defined by the condition $\phi_q(\beta=2,W_{\mathrm{MBL}}^{\mathrm{typ}})=0$. As previously discussed, the quenched free-energy suppresses the contribution of rare, large outliers, making it a good indicator of the typical behavior at a given system size L. Consequently, we expect this estimator to capture the localization transition of typical disorder realizations within the accessible sizes. It should therefore align with the critical disorder strength estimated through standard spectral observables and conventional approaches to the MBL transition. As will be discussed below, this may also apply to W_{ergo} . For instance, applying this analysis to the single-particle Anderson localization (AL) on random regular graphs (RRG), one finds that $W_{\mathrm{AL}}^{\mathrm{typ}}$ coincides with the disorder strength at which

the average gap ratio shows an apparent transition from RMT to Poisson behavior, with $W_{\rm ergo} \lesssim W_{\rm AL}^{\rm typ}$.

It is important to note that the phase diagram derived from these estimators is a finite-L phase diagram. We observe significant finite-size effects and drifts in the positions of the transitions between different regimes (see Figs. 3.14). Concerning the fate of this phase diagram in the $L \to \infty$ limit, two main scenarios are possible. In the first scenario, the intermediate phase $W_{\rm ergo} < W < W_{\rm MBL}$ —where delocalization occurs through rare resonances—is merely a finite-size crossover region that progressively shrinks and eventually disappears as $L \to \infty$, yielding a direct transition from the fully chaotic phase to the MBL phase (as for the Anderson model on the RRG). In the second scenario, this intermediate phase persists in the thermodynamic limit as a genuine new type of delocalized phase. Our numerical results for the accessible system sizes seem to favor the first scenario: the intermediate region shrinks progressively as L increases (see Sec. ??).

In the first scenario, the apparent non-ergodic behavior is a statistical artifact: typical finite-L samples are not representative of the thermodynamic limit. As established earlier, for $W > W_{\rm MBL}^{\rm typ}$, typical instances are localized with exponentially decaying long-range resonances, yet rare disorder realizations—identified by $W_{\rm MBL}(L)$ —can still drive delocalization. Although such realizations are uncommon at currently accessible sizes, they become typical at larger L. Therefore, in this scenario, it is natural to identify $W_{\rm MBL}^{\rm typ}$ as the disorder strength at which apparent ergodicity is effectively broken for finite systems. We thus expect $W_{\rm MBL}^{\rm typ}$ and $W_{\rm ergo}(L)$ to exhibit the same trend, both drifting toward $W_{\rm MBL}$ as $L \to \infty$. In the second (less likely) scenario, $W_{\rm MBL}^{\rm typ}$ would still drift toward $W_{\rm MBL}$ in the thermodynamic limit, while $W_{\rm ergo}$ would converge to a finite value strictly smaller than both $W_{\rm MBL}$ and $W_{\rm MBL}^{\rm typ}$.

In the rest of the article we systematically track the behavior of β_{\star} , $\phi_a(\beta_{\star})$, and $\phi_q(\beta=2)$, varying the parameters of the model. We present results in both the spin and Anderson bases, identified by different color scales in the following figures.

3.6 Finite size analysis

The free-energies and the relevant indicators

InFig. 3.12 we show examples of the free-energy densities as a function of the auxiliary parameter—or 'inverse temperature'— β . In the Figure we show the case $\Delta=1$ for several disorder strengths across the phase diagram, and for system sizes, shown in the color scale on the right of the figures. Results are shown for both spin basis (red, top row) and Anderson basis (blue, bottom row). The corresponding biased Hilbert space Landauer transmission, $\mathcal{T}_0(\beta)$, is computed using initial states $|0\rangle$ taken from the middle of the many-body spectrum. This choice, along with other simulation details, are described in App. A1.

The solid lines denote the free-energy functions $\tilde{\phi}_a$, i.e. the annealed free-energy with the replacement defined in Eq. (3.127). The increasing part of the annealed free-energy functions ϕ_a for $\beta > \beta_\star$ are shown as dashed lines, with their respective minima marked by vertical dashed lines of the same color. The error bars associated with these annealed free-energy functions are omitted for $\beta > \beta_\star$, as the average value of \mathcal{T}_0 becomes ill-defined in this regime, and ϕ_a depends strongly on the number of samples, being dominated by extreme value statistics.

The quenched free-energy functions $\phi_q(\beta)$ are also shown in solid lines with triangular markers. The two relevant values $\beta=2$ and $\tilde{\phi}_a=\phi_q=0$ are also shown as vertical and horizontal dashed-gray lines, respectively.

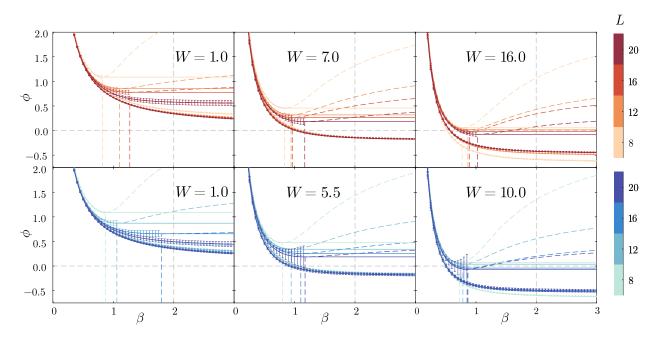


Fig. 3.12: The annealed free-energy ϕ_a (dashed), the modified annealed free-energy $\tilde{\phi}_a$ (solid) and the quenched free-energy ϕ_q (solid with triangular markers) in the spin (warm colors) and Anderson (cold colors) bases. Low (left panel), intermediate (middle panel) and large (right panel) disorder widths. The sizes are distinguished by the colors of the scale. The dashed gray lines show the relevant values at $\beta=2$ (physical transport) and $\phi(\beta)=0$ (delocalization/localization). The vertical colored dashed lines show the position of β_{\star} , for each curve.

For $\beta < \beta_\star$, the annealed free-energy $\phi_a(\beta)$ (and accordingly its modified counterpart $\tilde{\phi}_a(\beta)$) closely follows the quenched free-energy $\phi_q(\beta)$, with the gap between them narrowing as the system size L increases. In the vicinity of $\beta = \beta_\star$ the annealed free-energy $\phi_a(\beta)$ begins to deviate, developing into a convex function with a minimum. Both β_\star and the corresponding value of the annealed free-energy $\phi_a(\beta_\star)$ exhibit a systematic drift with system size, while the value of the quenched free-energy $\phi_q(\beta=2)$ shows weaker finite-size effect.

For small disorder (W=1, left panels), we find that $\beta_{\star}>2$ as the system size is increased, indicating that \mathcal{T}_0 receives contributions from an exponential number of terms. In contrast, for both intermediate (W=5.5 and W=7, middle panels) and strong disorders (W=10 and W=16, right panels), we observe $\beta_{\star}<2$, signaling that, at least for the accessible system sizes, the sum (3.114) is dominated by a few outliers in the tails of the distributions of the propagators.

For the intermediate disorder regime $\phi_a(\beta_\star)>0$, implying that the typical value of \mathcal{T}_0 in the large L limit should remain finite due to the inclusion of rare, large outliers of \mathcal{T}_0 . Upon increasing even more the disorder, however, $\phi_a(\beta_\star)<0$, indicating that the typical value of \mathcal{T}_0 vanishes exponentially upon increasing the size of the system. The results in both bases exhibit qualitatively similar behavior and allow us to distinguish the three relevant regimes. Yet, the crossover between these regimes occur at different values of disorder strength W, indicating a quantitative difference between the two bases—a point to which we will return in later sections.

3.6.1 The finite size phase diagrams

In what follows, we show the behaviour of β_{\star} , $\phi_a(\beta_{\star})$ and $\phi_q(\beta=2)$ as a function of W with increasing L. From these curves we extract the relevant finite-size critical disorder widths

 $W_{\rm ergo}(L)$, $W_{\rm MBL}(L)$ and $W_{\rm MBL}^{\rm typ}(L)$, respectively, as it is shown in the example of Fig. 3.13 for $\Delta=1.0$. From top to bottom the panels are: the minimum β_{\star} , the height of the plateau $\phi_a(\beta_{\star})$, and the value of the quenched free-energy for the physical transport $\phi_q(\beta=2)$. The critical disorder widths are computed through a linear or cubic spline interpolation using the points closest to the relevant crossings. The error bars associated to the critical disorder widths are computed through standard error propagation—whenever the corresponding covariance matrix is well-behaved, see App. A2 for further details—or with a Monte-Carlo based bootstrap resampling.

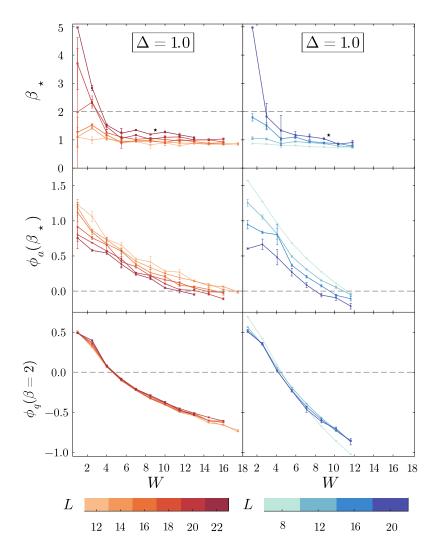


Fig. 3.13: (From top to bottom) Disorder width W dependence of β_{\star} , $\phi_a(\beta_{\star})$, and $\phi_q(\beta=2)$ for different system sizes, in the spin basis (left panels) and the Anderson basis (right panels). The relevant values $\beta_{\star} =$ 2 and $\phi = 0$ are correspondingly indicated with a horizontal dashed-gray line. The crossing of the curves with these lines identifies the position of $W_{\text{ergo}}(L)$, $W_{\mathrm{MBL}}(L)$ and $W_{\mathrm{MBL}}^{\mathrm{typ}}(L)$, accordingly. The star symbol at disorder W = 9 in the first column is obtained by computing the same quantity, β_{\star} , through another method, in order to verify the consistency of the results (see Sec. 3.6.3).

We repeat this procedure for several values of Δ to construct a finite size phase-diagram in the $W-\Delta$ plane. The results of this procedure are shown in Fig. 3.14 for both the spin and Anderson bases.

We identify three distinct finite-size regimes. At low disorder strengths, where $\beta_{\star}>2$ and $\phi_{q,a}(\beta=2)>0$, the system is in an *ergodic regime*, in which $\mathcal{T}_0(\beta=2)$ receives contributions from an exponentially large number of terms in the sum. As a result, it is self-averaging, its distribution has rapidly decaying tails, and sample-to-sample fluctuations are small. The typical and average values of $\mathcal{T}_0(2)$ both increase with the number of target states connected to the initial state $|0\rangle$, which grows exponentially with L. This behavior corresponds to a standard metallic regime for the conductance on the Hilbert space graph, meaning that the probability of delocalization from a random initial state $|0\rangle$ approaches 1 as L increases. This regime is depicted in green in the leftmost part of the phase diagrams Fig. 3.14.

A second regime emerges at intermediate disorder, $\beta_{\star} < 2$ and $\phi_a(\beta_{\star}) > 0$. This implies

that $\mathcal{T}_0(\beta=2)$ is dominated by rare outliers in the tails of the distribution of the propagators, corresponding to rare system-wide anomalously large resonances between distant configurations of the system. In this regime the distribution of $\mathcal{T}_0(2)$ develops heavy tails. According to our analogy with mean-field classical disordered systems, its typical value at large L will be ultimately dominated by samples that are rare at the accessible system sizes, and that feature anomalously large transmission events, leading to delocalization through rare resonances. As a result, the typical value of the probability that the system reaches one of the configurations at zero overlap from a random initial state will eventually approach 1 in the thermodynamic limit. However, only a few specific disorder-dependent configurations will be reached under the unitary evolution, corresponding to an extremely heterogeneous spreading of the wave-packet on the Hilbert space graph. This regime is depicted in green in Fig. 3.14, and is separated from the weak-disorder ergodic regime by the crossover line $W_{\text{ergo}}(L, \Delta)$, indicated by square markers (darker colors correspond to larger L).

Note that the line where $\phi_q(\beta=2)=0$, which defines the typical disorder strength $W_{\rm typ}$, lies within the intermediate region of the phase diagram. This line separates a regime at $W < W_{\rm typ}$, where typical samples are delocalized for the system sizes accessible numerically, from a regime of stronger disorder, where typical samples are localized. The position of this line agrees well with previous estimates of the MBL transition based on standard observables and conventional approaches [246, 247]. Importantly, most of the intermediate region corresponds to parameters where typical samples appear localized. This implies that rare resonances, ultimately responsible for delocalization to distant configurations over very long times in the asymptotic regime, are typically absent in the disorder realizations we can currently probe numerically. Nevertheless, our approach, inspired by mean-field theories of disordered glassy systems, provides a way to capture the asymptotic effects of such atypical disorder realizations.

At stronger disorder, the system enters a third regime, characterized by $\beta_{\star} < 2$ and $\phi_a(\beta_{\star}) < 0$. This corresponds to a genuine localized behaviour, as the typical value of the transmission $\mathcal{T}_0(\beta=2)$ vanishes (exponentially) in the large L limit even with the contributions of anomalously large outliers coming from the right tails of \mathcal{G}_{0f} . This phase is shown in shades of purple, with the corresponding crossover lines $W_{\text{MBL}}(L,\Delta)$ marked by circular symbols. The graded coloring within each region serves as a visual guide to the system sizes L used in the analysis: darker tones indicate larger system sizes.

A clear finite-size trend is observed: as L increases, $W_{\rm ergo}(L)$ shifts toward higher disorder (rightward), while $W_{\rm MBL}(L)$ shifts toward lower disorder (leftward), reflecting a systematic drift of the phase boundaries with increasing system size. This implies that the broad regime where the system delocalizes through rare events shrinks upon increasing L. This leaves open the possibility that the entire intermediate region corresponds to a finite-size prethermal crossover [257, 258], and that in the thermodynamic limit there is a direct transition from the ergodic phase to the MBL phase, i.e. that the two transition lines $W_{\rm ergo}(L,\Delta)$ and $W_{\rm MBL}(L,\Delta)$ may converge as $L\to\infty$.

This systematic drift of the transition point differs significantly from what was observed when applying the same method to the random field Ising model in a transverse field (also known as the Imbrie model), studied in Ref. [281]. The precise origin of the discrepancy between the two models remains unclear at present, though it certainly calls for further investigation.

Furthermore, the crossover lines found in the spin and in the Anderson basis are quantitatively different. This difference becomes more pronounced at small Δ . In particular the spin basis seems to perform poorly in detecting localization at small values of the interaction. This is especially evident at $\Delta=0$ —indicated by the horizontal black line—where the eigenstates of the Hamiltonian are tensor products of eigenstates of the single-particle Anderson model

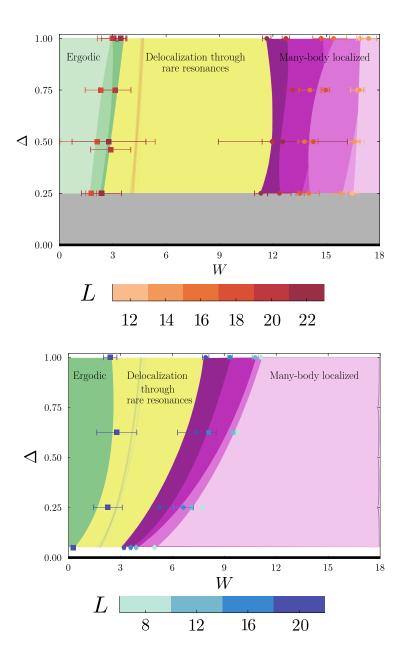


Fig. 3.14: Phase diagram of the system at the center of the energy spectrum, shown in the Δ –W plane. Crossover lines between the three regimes are indicated as follows: $W_{\rm ergo}$ (square markers), $W_{
m MBL}$ (circular markers), and $W_{
m MBL}^{
m typ}$ (solid transparent lines), each plotted for the system sizes indicated in the color scale. The three identified regimes—ergodic, delocalization via rare resonances, and many-body localization—are shaded in cyan, green, and purple, respectively. Gradients within each colored region serve as a visual guide to distinguish crossover lines estimated for different system sizes: regions identified from smaller sizes are darker and get lighter upon increasing L.

in one dimension. These single-particle eigenstates are known to be localized for any finite value of the disorder strength W, with a disorder-dependent characteristic localization length scale $\xi_{\rm loc}$. These localized eigenstates are not aligned with the spin basis states. As a result, when the system is initialized in a random spin basis state, the wavepacket spreads dynamically until it 'accommodates' into a superposition of eigenstates where the initial condition has a strong support. This evolution corresponds to a partial spreading—from a single spin basis state to a 'blob' of eigenstates whose typical size scales as $\xi_{\rm loc}^{L/2}$, in which the system is properly localized at sufficient strong disorder. However, to correctly recognize this partial spreading does not correspond to proper delocalization, it is necessary that the system size satisfies $L \gg \xi_{\rm loc}$.

For the system sizes accessible in numerical simulations, this condition is only met at sufficiently strong disorder, where the localization length is small enough compared to the system size. At smaller disorder, the localization length $\xi_{\rm loc}$ becomes large, introducing strong finite-size effects that hinder the observation of the localization behavior. This implies that at small Δ and moderate W, the spin basis is not well suited for the method, as it differs too much from the basis formed by the many-body localized eigenstates of the Hamiltonian.

This is why, in our analysis using the spin basis, the transition lines appear essentially vertical as Δ decreases. In particular, we can explicitly verify that our method in the spin basis fails to correctly capture localization at small W when $\Delta=0$. For these reasons, we have restricted our analysis in the spin basis to the regime $\Delta>0.25$, below which it becomes unreliable.

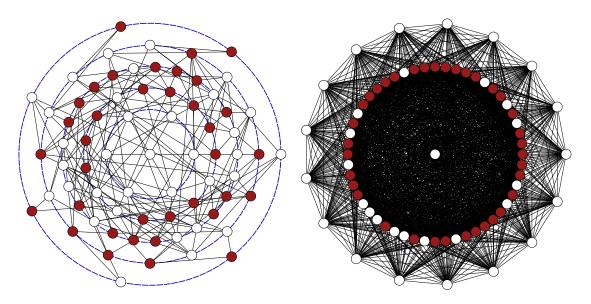


Fig. 3.15: Hilbert space graphs for L=8 shown in the spin basis (left) and the Anderson basis (right). The central vertex represents a random initial condition in the middle of the energy spectrum, that is a basis state of the Hamiltonian. All vertices of the graph are connected with black edges denoting the distance in the Hilbert space graph, given by number of applications of the Hamiltonian. In the spin basis we use concentric blue-dashed circles to aid the identification of these equidistant vertices, as the structure in this case is more irregular. The red colored vertices correspond to basis states belonging to the equator set $\mathcal E$ in both bases.

Additionally, the spin basis systematically predicts finite-size transitions at stronger disorder values compared to the Anderson basis. This leads to a substantial quantitative difference in the width of the regime where delocalization is driven by rare events. The origin of this discrepancy appears to be twofold.

First, as discussed above, the nature of the initial state plays a crucial role and depends strongly on the basis in which it is defined. In the Anderson basis, each basis state corresponds to a specific configuration of localized single-particle orbitals and is therefore a local modification of an eigenstate of the local integrals of motion (LIOMs). As such, an initial condition prepared in an Anderson basis state is expected to exhibit a small spreading in the MBL phase. In contrast, an initial state defined in the spin basis is expected to spread over many other vertices in Hilbert space—specifically, over all those LIOM eigenstates on which the initial state has significant projection—even in the MBL phase. This initial spreading leads to pronounced finite-size effects whenever $L \sim \xi_{loc}$. Secondly, there is a mismatch between the equator sets $A_0(q=0)$ and $S_0(q=0)$, used to identify target basis states at large distance from the initial state $|0\rangle$. In the Anderson basis, all target states $|f\rangle \in \mathcal{E} = \mathcal{A}_0(q=0)$ lie at the same graph distance from $|0\rangle$ on the Hilbert space graph, (defined as the length of the shortest path between $|0\rangle$ and $|f\rangle$, i.e., the minimum number of applications of the off-diagonal part of the Hamiltonian needed to connect them). In contrast, the target states in the spin basis, $|f\rangle \in \mathcal{E} = \mathcal{S}_0(q=0)$, have varying distances from $|0\rangle$, as different vertices have different degrees. The degree depends on the number

of domain walls present in each basis state, which can fluctuate from 2 to L, as explained in Sec. 3.3.2.

This difference between the two basis is illustrated in Fig. 3.15, where we present the complete Hilbert space graph for the case L=8 in both bases. The central vertex represents a randomly chosen initial condition with energy near the middle of the many-body spectrum. All edges connecting the Hilbert states through applications of the off-diagonal part of the Hamiltonian are shown. Vertices are arranged radially outward according to their distance on the graph from the central vertex $|0\rangle$. In the spin basis, we overlay concentric circles to help identifying the vertices that are equidistant from $|0\rangle$, as the structure in this case is more irregular due to the fluctuating connectivity the nodes. While in the Anderson basis all target states lie at the same graph distance from the initial condition, the targets in the spin basis are more dispersed: a large fraction of them are at a distance L/4 on the graph, but others are found at larger distances. Some of these target states are even located at the maximal possible distance from $|0\rangle$, reflecting the broader distribution of distances between spin configurations at zero overlap in the spin basis.

However, as the system size increases, the differences between the two computational bases are expected to diminish. On the one hand, due to the local nature of the LIOMs, when $L\gg\xi_{\rm loc}$, the 'blobs' representing LIOM eigenstates in the spin basis become effectively point-like: their typical spatial extent remains much smaller than the total Hilbert space volume. On the other hand, as L increases, most target states tend to concentrate uniformly around a distance L/4 from the initial state, since the majority of nodes on the Hilbert space graph contain approximately L/2 domain walls. As a result, we expect the transition lines to the MBL regime to eventually converge to a common value, independent of the choice of computational basis.

Remarkably, even very weak interactions (e.g., $\Delta=0.05$ in the Anderson basis, see Fig. 3.14) yield a finite typical value of \mathcal{T}_0 , signaling delocalization for the finite system sizes considered. In other words, both transition lines, $W_{\rm MBL}$ and $W_{\rm MBL}^{\rm typ}$, remain finite at small Δ , indicating a discontinuous departure from the Anderson insulator at $\Delta=0$, where localization persists at arbitrarily small disorder W. This provides further numerical evidence of the non-perturbative effect of interactions, consistent with the spin–spin correlation analysis and the updated XXZ phase diagram reported in Refs. [279, 280].

3.6.2 Dependence on the target states distance

In this section, we modify the selection of target states, which so far has been restricted to the equator set \mathcal{E} . Specifically, to investigate the progressive spreading of the wave packet onto configurations at increasing distances, we vary the parameter q in the sets $\mathcal{S}_0(q)$ and $\mathcal{A}_0(q)$, defined in Eq. (3.107), for the spin and Anderson bases, respectively. We reintroduce the superscript on the overlap value, $q \to q^{S,A}$, to explicitly indicate the basis—spin (S) or Anderson (A)—in which the overlap is being measured. We reparametrize these overlaps to define a *correlation distance* $\zeta^{S,A}$, which serves as an ultrametric distance in Hilbert space, and is defined by

$$\zeta^{S,A} = 1 - q^{S,A} \ . \tag{3.129}$$

where $\zeta^{S,A}=0$ trivially corresponds to the initial random basis state itself, while $\zeta^{S,A}=1$ represents completely uncorrelated states, where half of the spins have been flipped in terms of the spin basis or, equivalently, where half of the spinless-fermions have hopped to different orbital occupations, in the Anderson basis. We will omit the cases $1<\zeta^{S,A}<2$, that correspond to anti-correlated basis states with respect to the initial condition $|0\rangle$. We refer to $\zeta^{S,A}$ as the *correlation distance*. In what follows we will restrict the analysis to $\Delta=1$.

We have computed $\mathcal{T}_0(\beta)$ while varying the correlation distance $\zeta^{S,A}$ between the target nodes and the initial condition, for several values of the disorder strength W and system sizes L, averaging over many disorder realizations and many initial states $|0\rangle$ randomly chosen close to the middle of the many-body spectrum. The procedure follows the same steps as in previous sections, with the only difference being the new dependence of the target states on $\zeta^{S,A}$. We evaluate the annealed free-energy $\phi_a(\beta,L,\zeta^{S,A})$ and extract the corresponding values of β_\star and $\phi_a(\beta_\star)$ for different values of the disorder close to the ergodic and MBL transitions, respectively. From these, we determine the characteristic distances that fulfill $\beta_\star(\zeta^{S,A})=2$ and $\phi_a(\beta_\star,\zeta^{S,A})=0$ by performing linear interpolations. An example of this procedure is shown in Fig. 3.16 for the spin basis, with L=20 and $\Delta=1$, an equivalent example for the Anderson basis is shown in App. A4.

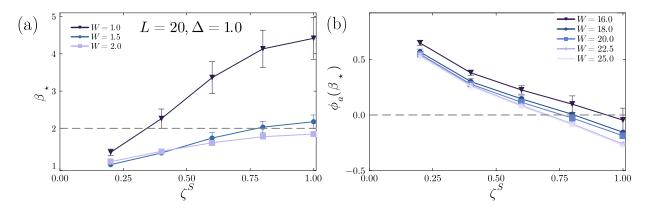


Fig. 3.16: Calculation of (a) β_{\star} and (b) $\phi_a(\beta_{\star})$ as functions of the correlation distance ζ^S , for $\Delta=1$ and L=20. The values of the disorder widths considered are shown in the legend. Horizontal gray dashed lines indicate the reference values $\beta_{\star}=2$ and $\phi_a(\beta_{\star})=0$.

These characteristic distances define crossover lines that separate different regimes in the W- ζ plane, as shown in Fig. 3.17 for both the spin (left) and Anderson (right) bases. For the condition $\beta_{\star}(\zeta^{S,A})=2$, we used the largest system size, L=20, where such a value of $\zeta^{S,A}$ is obtained. In contrast, for the condition $\phi_a(\beta_{\star},\zeta^{S,A})=0$, we display the results for three different system sizes: L=12,16, and 20.

In the first region (shaded green), the system remains ergodic within the distance defined by $\beta_{\star}(\zeta^{S,A})=2$: A randomly chosen initial basis state has a high probability of spreading uniformly to any other state in Hilbert space within the corresponding correlation distance. Interestingly, the curve $\beta_{\star}(\zeta^{S,A})=2$ bends to the right (i.e., toward larger disorder values) as the distance increases. This indicates the existence of a disorder window $(1.5 \lesssim W \lesssim 3)$ approximately) where the spreading of the wave packet is inhomogeneous and driven by rare resonances at short distances, but recovers a uniform, ergodic-like spreading over an exponential number of configurations at larger distances. A similar behavior is observed on the metallic side of the Anderson model on the Bethe lattice [350, 419].

In the white region, the spreading of the wave packet from the initial state $|0\rangle$ is highly inhomogeneous and dominated by a few rare resonances in the broad tails of the propagator distribution at the corresponding distances. As explained above, at weak disorder, for $W < W_{\rm ergo}$, the system eventually recover an ergodic behavior at large distances, after crossing the crossover line separating the white from the green region. In the intermediate regime, $W_{\rm ergo} < W < W_{\rm MBL}$, the dynamics is still driven by rare resonances up to the distances corresponding to states with zero overlap with $|0\rangle$. As a result, the spreading remains inhomogeneous across the entire Hilbert space. For $W > W_{\rm MBL}$, the transport is dominated by a few O(1) resonances at short distances; beyond that, i.e., for distances within the gray regions, they become so rare that the probability for a random initial state

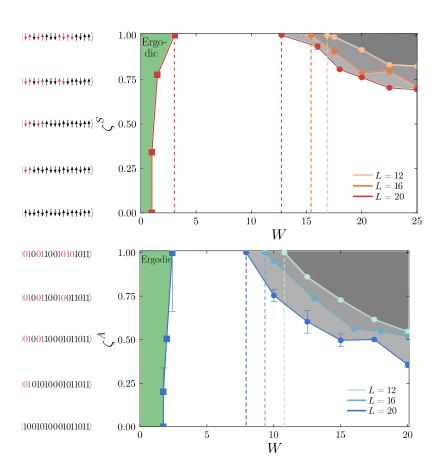


Fig. 3.17: Relevant regions in the $W-\zeta^{S,A}$ plane for bosons (top) and fermions (bottom) at $\Delta = 1$. The transition to the ergodic region (green) is tained from $\beta_{\star}(\zeta^{S,A}) = 2$, while the boundaries of the inaccessible regions follow (grav shades) $\phi_a(\beta_{\star}, \zeta^{S,A}) = 0$ for each system size. Critical disorder widths $W_{\rm ergo}$ and $W_{\rm MBL}$ are shown as dashed lines, colored by system size. On each vertical axis, a random initial basis state ($\zeta^{S,A} = 0$) is indicated, followed by example basis states at increasing distances $\zeta^{S,A}$. Equator states at $\zeta^S = 1$ represent complete decorrelation from the initial state.

to reach such distant configurations decays exponentially with system size—even when accounting for the statistical weight of rare events. The crossover distance, determined by the condition $\phi(\beta_\star,\zeta^{S,A})=0$ and marking the boundary of the gray regions, represents the maximal Hilbert space correlation length accessible under unitary dynamics from a typical initial state. Consequently, in the large-L limit, the gray regions become asymptotically inaccessible for typical initial conditions and disorder realizations.

As the disorder strength grows, the system becomes progressively confined to a smaller portion of Hilbert space near the initially prepared configuration. In other words, resonant transmissions become increasingly short-ranged. This behavior becomes more pronounced with larger system sizes, as the inaccessible region of Hilbert space expands, restricting more and more the set of basis states available for delocalization. This behavior reflects the persistent memory of the initial condition at strong disorder, where the values along the transition line $\phi_a(\beta_\star)=0$ define the maximum correlation distance (i.e., the minimal overlap) the system can explore.

Note that even deep within the MBL phase (e.g. W=20) there still exist rare resonances in the Hilbert space that allows the system to flip a finite fraction of the spins and partially decorrelate from the initial condition. This is consistent with the picture of resonant cat states [438], formed by two nearly degenerate spin basis states that differ by a fraction of spin flips. These rare events can drive partial delocalization even in strongly disordered regimes, within the system sizes considered.

3.6.3 The distributions of the Hilbert space Landauer tranmssions

In this section, we will explicitly examine the probability distributions of the Hilbert space Landauer transmissions \mathcal{T}_0 , for several disorder widths W and system sizes L. In Fig. 3.18, we show the results for three different system sizes (L=12, 16, 20) and three disorder strengths (W=1, 9, 20).

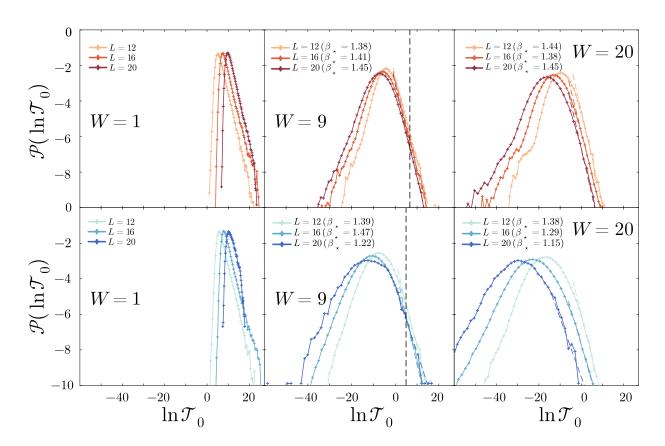


Fig. 3.18: Probability distributions for the Hilbert space Landauer transmission \mathcal{T}_0 , for the different sizes displayed in the key. For small (left), mid (center), and strong (right) disorder widths, in both spin (top panels) and Anderson (bottom panels) bases. These plots are the counterpart of the probability distributions of Fig. 3.9, where the probability to delocalize to the states at zero overlap from the random initial state is replaced by our proxy \mathcal{T}_0 . The vertical black-dashed line in the middle panel corresponds to the Hilbert space Landauer transmission, from a disorder realization that give rise to cat states with system-wide resonances, for L=20 (see the main text for a detailed explanation).

At weak disorder, W=1, the peak of the distribution shifts rightward with increasing L, and the distributions develop a sharper cutoff. This behavior indicates the absence of rare outliers and suggests that the typical and average values of \mathcal{T}_0 are proportional to each other, growing as a power of L due to the increasing number of outgoing channels $\mathcal{N}_{\mathcal{E}}$. This reflects the fact that $\phi_{q,a}(\beta=2)>0$, signaling an ergodic regime in which the typical value of \mathcal{T}_0 receives contributions from an exponential number of target states $|f\rangle$.

In contrast, for stronger disorder values, W=9 and W=20, the peak of the distribution shifts leftward, while the tails remain broad. This indicates the presence of significant fluctuations, causing the average and typical values of \mathcal{T}_0 to differ substantially. According to Derrida's theory of the freezing transition of directed polymers and its generalizations [421, 439] the exponent governing the tail of the probability distribution of $\mathcal{T}_0(\beta)$ is related to the

freezing inverse temperature β_{\star} via:

$$\mathcal{P}(\mathcal{T}_0(\beta)) \simeq \frac{e^{-\beta_\star^2 \ln \mathcal{N}_{\mathcal{E}} \phi_q / \beta}}{\mathcal{T}_0^{1+\beta_\star / \beta}} , \qquad (3.130)$$

This behavior has a clear intuitive origin: for $\beta < \beta_{\star}$, the typical and average values of \mathcal{T}_0 remain proportional, since the average is dominated by the bulk of the distribution rather than its tail. In contrast, for $\beta > \beta_{\star}$, the average becomes dominated by rare, large fluctuations, as the tail of the distribution decays with an exponent smaller than 2. We test this prediction by extracting the value of β_{\star} from the power-law fit of the tails of the distribution $\mathcal{P}(\mathcal{T}_0(\beta=2))$. The values of β_{\star} found from the fits of the power-law tails of the distribution are displayed in the key of Fig. 3.18 and shown in Fig. 3.13 with a star symbol for the largest system size L=20 and disorder width W=9, showing consistency with the value of β_{\star} found from the position of the minimum of the annealed free-energy.

Furthermore, in a recent study by some of the co-authors [438], the existence of an intermediate, non-ergodic phase in the disordered Heisenberg chain was linked to the emergence of unusual high-energy eigenstates exhibiting anomalously strong longitudinal spin–spin correlations [279, 280]. These eigenstates appear in nearly degenerate pairs, sparsely distributed across the exponentially large many-body spectrum. Remarkably, their properties are accurately captured by a simple toy model of *cat states*. These cat states take the form $|\psi\rangle_+ \sim |I_1\rangle \pm |I_2\rangle$, where $|I_1\rangle$ and $|I_2\rangle$ are spin-basis states.

The occurrence of such cat eigenstates is frequent in the intermediate disorder regime ($W \sim 10$), but they become increasingly rare at stronger disorder ($W \gtrsim 20$). Intuitively, the presence of a resonant cat state in a localized system can enhance the probability of delocalization: if a random initial spin basis state has a significant projection with one of these cat states, a resonance between $|I_1\rangle$ and $|I_2\rangle$ enables the system to explore both configurations, thus promoting delocalization. However, such states are rare in the spectrum, and their impact on the decorrelation from a typical random initial state—and their direct connection to delocalization via rare resonances—remains unestablished.

To fill this gap, we test this picture with our observables defined in Hilbert space, and we compute the Hilbert space Landauer transmission \mathcal{T}_0 —averaged over several initial conditions—for a specific disorder realization at system size L=20 and disorder strength W=9. This particular realization hosts pairs of nearly degenerate eigenstates exhibiting strong spin–spin correlations, consistent with the cat-state scenario [438]. The corresponding value of \mathcal{T}_0 is indicated by the black dashed line in the middle panel of Fig. 3.18 for W=9. Interestingly, this sample is also classified as a rare event in terms of the Hilbert space Landauer transmission, as its corresponding \mathcal{T}_0 lies within the tail of the distribution, where $P(\mathcal{T}_0(\beta=2))\approx 10^{-6}$. This observation suggests a potential link between atypical values of Hilbert space observables—such as \mathcal{T}_0 —and real space features like longitudinal spin–spin correlations. This intriguing connection merits further investigation.

3.7 Rare vs typical samples

One of the key features of our proxy observable for the probability to decorrelate from a random initial state, \mathcal{T}_0 , is that it naturally allows us to distinguish between typical and rare samples. In this section, we take advantage of this property to explore how the spectral and transport features differ when the quenched disorder corresponds to rare versus typical realizations. Specifically, rare samples are defined as disorder realizations of the random fields h_i for which the corresponding value \mathcal{T}_0 , averaged over multiple initial conditions, lies in the tails of the distributions shown in Fig. 3.18. In contrast, typical samples are

selected from the vicinity of the peak of the \mathcal{T}_0 distribution, representing the most probable values. We denote these rare and typical realizations of random fields as $\{h_i\}_{\text{rare}}$ and $\{h_i\}_{\text{typ}}$, respectively.

The structure of the eigenstates' amplitudes

We first probe these typical and rare samples by diagonalizing their associated Hamiltonian, through an implementation of Chebyshev filter diagonalization [440] that finds eigenpairs within the center of the many-body spectrum [441]. We perform this procedure for small (W=1), intermediate (W=9) and strong (W=20) disorders, for a fixed system size L=20.

For each of the several eigenstates found—here numbered with the sub-index α —we extract the basis state with largest wavefunction amplitude, i.e. $I \equiv \arg\max_{I'} |\psi_{\alpha}(I')|^2$. For a given eigenstate α , we order the amplitudes of all other basis states $|\psi_{\alpha}(J)|^2_{J\neq I}$, in terms of the correlation distance to the most probable state I. This correlation distance is given by

$$\zeta_{IJ}^{S,A} = 1 - q_{IJ}^{S,A} \,, \tag{3.131}$$

for the spin and Anderson bases, respectively. The basic idea is to study the decay of many-body eigenstates from their main peak, by averaging the wavefunction amplitudes of all states $\{|J\rangle\}$ at equal correlation distances ζ_{IJ}^S . The results of this procedure are shown in the top panel of Fig. 3.19, where we have calculated the corresponding eigenstates for both rare, $\{h_i\}_{\text{rare}}$, and typical, $\{h_i\}_{\text{typ}}$, realizations of the random fields. We will present the results associated with the Anderson basis in App. A4.

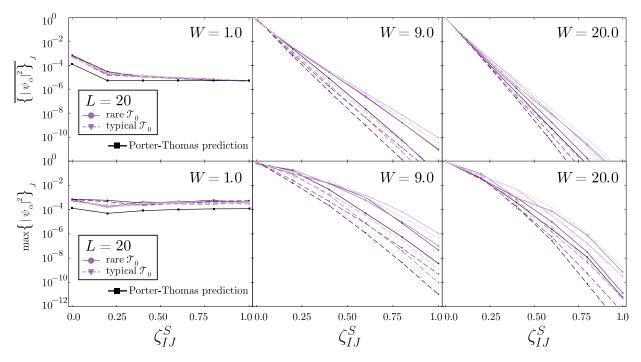


Fig. 3.19: Basis state amplitudes, within a given eigenstate, as a function of the correlation distance ζ_{IJ}^S from the basis state maximal amplitude. The top row of panels shows the average amplitudes over all spin states that share the same distance ζ^S , while the bottom panel displays the maximum amplitudes among those same basis states. Different color tones correspond to five distinct eigenstates. Dashed lines with triangular markers represent typical disorder realizations, whereas solid lines with circular markers denote rare realizations.

At weak disorder (W=1), the amplitudes of the eigenstates remain nearly constant across the Hilbert space graph, showing no decay with distance. Fluctuations between eigenstates are minimal, even across different disorder realizations. In contrast, at stronger disorder (W=9 and W=20), the eigenstates exhibit exponentially fast decay with the correlation distance $\zeta^{(S)}$, becoming strongly peaked around the reference basis state I with amplitudes of order 1. In this regime, differences between disorder realizations become apparent: the fluctuations in eigenstate amplitudes are meaningful—specially for the the intermediate disorder width of W=9—and eigenstates from rare disorder realizations tend to decay more slowly, by 2 orders of magnitudes close to $\zeta^{(S)}=1$, compared to those from typical ones.

To further explore this behavior, we modify the previous analysis by taking the maximum amplitude—rather than the average—among all basis states at a given correlation distance ζ^S . This modified approach is illustrated in the bottom panel of Fig. 3.19. The basic idea is to investigate whether rare, disorder-dependent resonances occurring at specific points in the Hilbert space graph lead to an inhomogeneous decay of the wave functions along different paths, resulting in a strong disparity between the average decay and the decay along the path corresponding to the maximal amplitude.

For weak disorder (W=1), the change is minimal: the decay remains largely unaffected, as the eigenstates display ergodic behavior. In this regime, amplitude fluctuations are small, and the wavefunction remains nearly uniform across the entire Hilbert space graph, equivalent to its averaged counterpart. In contrast, at strong disorder (W=9 and W=20), the use of the maximum amplitude reveals significant fluctuations between eigenstates, particularly between typical and rare disorder realizations. In these cases, eigenstates from typical realizations decay faster than those from rare ones. Moreover, examining individual eigenstates shows that the decay away from the maximum amplitude is highly anisotropic across the Hilbert space graph: Specific directions, aligned with the largest amplitudes exhibit a much slower decay (by 3 or 4 ourders of magnitudes close to $\zeta^S=1$, than the average one. In other words, there exist spin basis states at large Hilbert space distances whose amplitudes are anomalously large compared to the typical amplitudes at the same distance.

This effect is especially pronounced at intermediate disorder. It suggests that, in this regime, delocalization proceeds in a highly heterogeneous manner, along rare, disorder-dependent paths through the Hilbert space graph. These paths are determined by the presence of long-range resonances that connect distant basis states and dominate the eigenstate structure in certain realizations. We will further investigate this structure by explicitly examining the presence of resonant paths in the Hilbert space graph, as generated by both typical and rare realizations of the random disorder fields.

Rarefaction of paths on the Hilbert space graph

Our results so far suggest that there is a broad intermediate regime in which delocalization occurs via a small number of rare long-range resonances on the Hilbert space graph. This interpretation is also supported by the analysis of the eigenstate structure presented above. In this section, we further explore this scenario by directly investigating the paths on the Hilbert space graph that contribute most significantly to the delocalization probability from a random initial configuration. We adopt an approach that has previously been used by Lemarié in Ref. [289] to study the zero temperature properties of single-particle Anderson localization in a two dimensional geometry.

In the standard setup of quantum transport involving a scattering geometry, as illustrated schematically in Fig. 3.10, electrons are injected from the leads on the left and extracted

through those on the right. In the presence of strong disorder within the scattering region, electron transport becomes highly inhomogeneous. Rather than spreading uniformly, an electron at zero temperature follows a narrow, meandering path through the disordered potential landscape—effectively forming a 'trajectory' or conducting channel connecting the leads [413]. This behavior contrasts with the weak-disorder, diffusive regime, in which the electron's probability distribution is approximately uniform across the sample. A central challenge, therefore, is how to visualize or reconstruct these hidden transmission paths in the localized regime.

As introduced in Refs. [287, 288], one can devise a clever numerical perturbation technique, inspired by experimental scanning gate microscopy methods [442, 443] to visualize these dominant paths. In such experiments, a movable tip locally modifies the potential landscape of a nanoscale conductor, and the resulting changes in conductance reveal the regions through which current flows. The numerical analog operates as follows: for a given disordered sample, one slightly perturbs the on-site disorder at a specific location and measures how much the conductance g_0 between the incoming and outgoing leads changes. If the perturbed site lies along a main transmission path, even a local modification will significantly affect the coherent transport, resulting in a noticeable change in g_0 . Conversely, if the site is far from the dominant path (i.e., weakly visited by the electron's wavefunction), the conductance remains essentially unchanged. By systematically applying this 'poke test' across all sites, one obtains a spatial map of conductance sensitivity: regions where g_0 is highly responsive to local perturbations directly identify the dominant current-carrying pathways.

Lemarié's work builds on the same principle, applied to quantum transport of non-interacting electrons at zero temperature in 2d. In this approach, the on-site disorder potential ε_i of the single-particle Anderson model is locally perturbed according to $\varepsilon_i \to -\varepsilon_i$. The resulting conductance g_i is then computed for each perturbed site. To quantify the impact of the local perturbation, the relative conductance response at site i is defined as

$$\delta g(i) \equiv \frac{|g_i - g_0|}{q_0} \tag{3.132}$$

where g_0 is the original conductance of the sample, and g_i is the conductance after flipping the disorder at site i.

As in Anderson localization, we argue that in the strong disorder regime, the propagation of a fictitious particle in Hilbert space becomes highly inhomogeneous, dominated by an O(1) number of transmission paths on the Hilbert space graph. To probe this structure, we adopt an analogous approach to the one outlined above. To this aim, we formally reintroduce the semi-infinite leads: an incoming lead is connected to the initial basis state $|0\rangle$ —with energy in the middle of the many-body spectrum—and several outgoing leads connected to each of the target states $|f\rangle \in \mathcal{E}$. The fictitious particle is then injected at $|0\rangle$ and may propagate through the network of allowed transitions in Hilbert space before being absorbed at one of the target states. This defines an effective transport setup in Hilbert space, where the network of paths connecting $|0\rangle$ to the various $|f\rangle \in \mathcal{E}$ plays the role of the scattering geometry. For simplicity we only perform this analysis in the spin basis.

The effect of the leads is incorporated by adding a self-energy term, $-i\Sigma_I$, to the diagonal element of the Hamiltonian in Eq. (3.78) for the respective basis state at which the lead is connected. This is implemented via the replacement

$$\hat{\mathcal{H}} \to \hat{\mathcal{H}} - i\Sigma_I |I\rangle \langle I|$$
 (3.133)

with I corresponding to both the inital basis state $|0\rangle$, and its respective targets $|f\rangle \in \mathcal{E}$. The addition of an imaginary part of the self-energies of the leads is equivalent to inserting

dissipation to the Hilbert space-graph at the vertices of interest. A pictorial representation of this construction can be seen in Fig. 3.20.

In our case the local perturbations to the disordered potential are introduced indirectly, by modifying the diagonal elements of the sparse Hamiltonian, which encodes the quenched disorder from a specific realization of the random fields. Specifically, we alter the diagonal element associated with the basis state $|I\rangle$ —a vertex in the Hilbert space graph—through:

$$\mathcal{H}'_{II} = \mathcal{H}_{II} + 2(\overline{E} - \mathcal{H}_{II}) |I\rangle \langle I| , \qquad (3.134)$$

where \overline{E} is the average energy over all basis states—or equivalently, the average of the diagonal elements of the Hamiltonian matrix \mathcal{H} . In other words, this perturbation flips the value of the energy associated to the state $|I\rangle$ with respect to the mean \overline{E} . We repeat this procedure for each of the \mathcal{N} diagonal elements, and calculate the associated resolvent matrix each time. Instead of doing this exhaustively, we calculate it using perturbation theory, which becomes exact in this case and reduces to the Sherman-Morrison formula for inverting matrices under 1-rank perturbations. Using this formula

$$\mathcal{G}' = \mathcal{G} - \frac{2(\mathcal{H}_{II} - \overline{E}) \mathcal{G} |I\rangle \langle I| \mathcal{G}}{1 + 2(\mathcal{H}_{II} - \overline{E}) \mathcal{G}_{II}}, \qquad (3.135)$$

we can recompute the Hilbert space Landauer transmissions perturbed at each vertex I, which we call $\mathcal{T}_0^{(p)}(I)$, and we measure the (normalized) response defined as

$$\delta g(I) \equiv \frac{|\mathcal{T}_0^{(p)}(I) - \mathcal{T}_0|}{\sum_{I} |\mathcal{T}_0^{(p)}(I) - \mathcal{T}_0|}.$$
 (3.136)

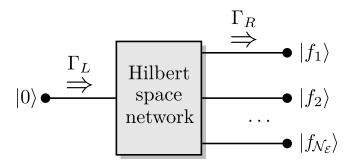


Fig. 3.20: Scattering geometry to measure the reaction of the Hilbert space conductance and responses to small perturbations. The initial basis state at the center of the many-body spectrum is connected to a semi-infinite lead through which fictitious particles are injected. Similarly, semi-infinite right leads are connected to the target vertices $|f\rangle$, that belong to the equator \mathcal{E} .

We have calculated the response δg at low, intermediate, and large disorder strengths. We have repeated the calculations for typical samples (for which \mathcal{T}_0 is in the bulk of the probability distribution) and for rare samples (corresponding to disorder realizations that produce values of \mathcal{T}_0 in the tails of the distribution). In order to visualize the effect of the perturbation—that can be interpreted as the conductance generated by the incoming

fictitious particles in the Hilbert space network—we normalize δg in the following way:

$$\delta g'(I) = \begin{cases} 0.5 \frac{\delta g(I) - \delta g_{\min}}{\delta g(0) - \delta g_{\min}}, & \text{if } \delta g(I) < \delta g(0) \\ 0.5 + 0.5 \frac{\delta g(I) - \delta g(0)}{\delta g_{\max} - \delta g(0)}, & \text{if } \delta g(I) > \delta g(0) \\ 0.5, & \text{otherwise.} \end{cases}$$

$$(3.137)$$

In words, we normalize the response $\delta g(I)$ onto a scale from 0 to 1. The initial basis state $|0\rangle$ (to which the incoming semi-infinite lead is connected) defines the reference of this scale, because we are interested in probing how strong or weak each other vertex's response is relative to the initial signal at $|0\rangle$ here denoted as $\delta g(0)$.

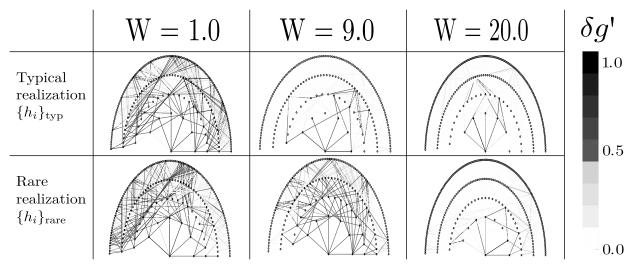


Fig. 3.21: Rarefaction of paths in the Hilbert space graph (for the spin basis) for L=16, for three disorder strengths. The color scale is set according to the central vertex $\delta q'=0.5$.

We show the results of this construction in Fig. 3.21 for L=16. The initial state $|0\rangle$ is placed at the center in each of the six diagrams, and all other vertices connected to it are also shown. These vertices are arranged radially outward, ordered according to their Hilbert space distance from $|0\rangle$, i.e. same radius implies the same distance in the graph from $|0\rangle$. The vertices are then colored according to the scale of $\delta g'$. Darker colors correspond to $\delta g'>0.5$, all the way to black for $\delta g'=1$. These are vertices that exhibit a strong response to the perturbation of \mathcal{T}_0 , indicating that the corresponding basis states lie along dominant transmission paths originating from the initial state $|0\rangle$. In contrast, vertices shown in lighter colors (with $\delta g'<0.5$) exhibit a weaker response, suggesting a low probability for the fictitious particle to reach those sites from the initial condition. Although the reaction $\delta g'$ is a variable assigned to the vertices of the graph, we have colored the edges to aide visualization. Hence, the edges are colored according to the value of the preceding vertex connected to them.

For weak disorder strength (W=1), both rare and typical disorder realizations exhibit a proliferation of transmission paths in Hilbert space, enabling the transport of the injected fictitious particles to distant vertices in a uniform wa, with most of the target states at large distance from $|0\rangle$ reached by dark paths. When disorder is increased (W=9), these reactions become smaller and dark conducting paths are strongly rarefied. For typical samples these rarefied conducting paths do not reach far away spin configurations, corresponding to the suppression of wave-packet spreading and localization. Yet, rare disorder realizations from

the tails of the distribution of \mathcal{T}_0 exhibit much more reactive paths than typical samples, some of them extending to a few basis states at large distances. This corresponds to a strongly inhomogeneous spreading of the wave-packet starting at $|0\rangle$, and delocalization along specific disorder-dependent paths, occurring only for rare disorder realizations. At very strong disorder W=20, finding subgraphs with highly-reactive paths is even more rare—although possible—but even these rare paths are unable to reach distant spin configurations, even for rare disorder realizations, corresponding to complete suppression of long-distance spreading of the wave-packet and a hallmark of genuine MBL.

More quantitatively, the contribution of these strongly reactive paths can be characterized using the (average and typical) inverse participation ratio (IPR) associated with the reaction amplitudes, they are defined as (recall that according to our definition (3.136) the response is normalized to one, i.e. $\sum_{I} \delta g(I) = 1$):

$$\mathcal{I}_{2}(\delta g) = \mathbb{E}\left[\sum_{I} \delta g(I)^{2}\right] ,$$

$$\mathcal{I}_{2}^{\text{typ}}(\delta g) = \exp\left(\mathbb{E}\left[\log \sum_{I} \delta g(I)^{2}\right]\right) ,$$
(3.138)

where $\mathbb{E}[\cdots]$ is taken for several initial conditions and disorder realizations. In Fig. 3.22, we show the behaviour of the IPR as a function of the system size L. The definition of the typical IPR suppresses the influence of rare events, thereby capturing the behavior of typical disorder realizations.

For weak disorder (W=1), the response values are approximately uniform across all $\mathcal N$ basis states. In this regime, the inverse participation ratio (IPR) scales as $1/\mathcal N$, decreasing exponentially with system size L and reflecting the contribution of an extensive number of transmission paths. In contrast, for stronger disorder (W=9,W=20), only a small O(1) number of sites—not growing with the size of the Hilbert space—exhibit significant responses, leading to an IPR that saturates to a finite plateau over the system sizes studied.

At small disorder (W=1), the typical and average IPR are essentially equivalent. At large disorder ($W=9,\,W=20$), the typical IPR is consistently slightly smaller than the average one. This small difference highlights the presence of rare disorder realizations with exceptionally reactive dominant paths, which significantly affect the arithmetic average.

These results are consistent with the picture developed in Sec. 3.7, which highlights the strong fluctuations in the structure of eigenstates near the middle of the spectrum. In the present context, this heterogeneity manifests as a broad distribution of dominant transmission paths in the Hilbert space graph, particularly pronounced in the intermediate disorder regime (W = 9).

3.8 Conclusions and outlook

In this work we have developed an approach based on an analogy with a class of mean-field disordered glassy systems that allows one to take into account the statistical weights of rare events for the MBL transition. We have extended the analysis of Ref. [281] by applying this method to the XXZ model varying the strength of the interaction Δ , considering both the spin and Anderson bases. Our study emphasizes the role of rare long-range resonances—arising in rare disorder realizations—in destabilizing the MBL transition at finite sizes in a broad intermediate disorder range. These rare system-wide resonances are identified as the outliers in the probability distribution of transition amplitudes between distant configurations of the system in Hilbert space, expressed in the chosen computational basis. Concretely, we use

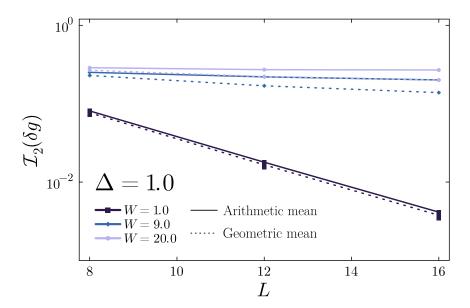


Fig. 3.22: Inverse participation ratios (\mathcal{I}_2) associated to the reaction term δg as a function of the system size L. The different curves correspond to average and typical IPR as in Eq. (3.138).

the propagators $|\mathcal{G}_{0f}|^2$ as a proxy for the probability that a system initialized in a random configuration $|0\rangle$ at time t=0 is found in the configuration $|f\rangle$ —located at large distance from $|0\rangle$ —at infinite time, $\sum_n |\langle f|n\rangle\langle n|0\rangle|^2$. Our study complements previous studies focused on real space observables and spectral signatures of rare events [251, 257, 258, 266, 270, 278–280], by providing a Hilbert space-based perspective.

To correctly evaluate the statistical weight of rare resonances in the asymptotic limit of large system size L, we exploited an analogy with classical disordered systems and introduced an auxiliary parameter β that plays the role of an effective temperature. This extension of the parameter space allows us to tune the influence of extreme outliers in the heavy-tailed distribution of propagators, and to identify, for each given disorder strength W, the value of β at which rare events begin to dominate the statistical measure.

This method reveals the existence of three distinct regimes: (i) an ergodic phase, (ii) an intermediate regime in which delocalization is driven by rare long-range resonances in an otherwise localized background, and (iii) a genuinely many-body localized phase, which remains stable even in the presence of anomalously large outliers emerging from the tails of the distributions of the propagators. We show that typical samples that we can probe numerically lack the system-wide resonances that ultimately lead to delocalization in the asymptotic limit in the intermediate regime. Yet, our approach inspired by the analogy with mean-field glassy systems captures their asymptotic contribution.

It is important to acknowledge the potential limitations and drawbacks of our approach. As explained in Sec. 3.5.2, for numerical convenience we do not perform a systematic study of the statistics of the 'true' order parameter for delocalization—namely, the typical value of the probability to delocalize from a random initial state after infinite time, $\mathbb{P}_{\mathcal{E}}$ —but instead study the typical value of the unregularized Hilbert space transmission \mathcal{T}_0 . This approximation is partially justified by the fact that the typical values of \mathcal{T}_0 and $\mathbb{P}_{\mathcal{E}}$ exhibit the same scaling with L when both decay exponentially with L (see Sec. 3.5.2 and Fig. 3.11), since in this regime the regularization of poles becomes asymptotically unnecessary. However, at finite L, the typical values of $\mathbb{P}_{\mathcal{E}}$ and \mathcal{T}_0 —obtained by computing $\mathbb{E}[\ln \mathbb{P}_{\mathcal{E}}]$ and $\mathbb{E}[\ln \mathcal{T}_0]$ —begin to decrease exponentially with L well before the MBL transition. In fact, throughout almost

the entire intermediate region (ii), which is dominated by rare resonances, both $e^{\mathbb{E}[\ln \mathbb{P}_{\mathcal{E}}]}$ and $e^{\mathbb{E}[\ln \mathcal{T}_0]}$ decay exponentially with L. We then employ a large-deviation approach to determine whether rare events in the tail of the probability distribution might alter this scaling when properly accounted for at large L. We find that in this intermediate regime, the system is indeed delocalized despite the exponential decay of $e^{\mathbb{E}[\ln \mathbb{P}_{\mathcal{E}}]}$ and $e^{\mathbb{E}[\ln \mathcal{T}_0]}$ at accessible system sizes.

A legitimate question then arises: could the large outliers that we suggest destabilize localization be overestimated due to the lack of regularization in \mathcal{T}_0 ? In this sense, our MBL threshold should be viewed as an upper bound—some of the rare events we consider crucial for MBL destabilization might be spurious artifacts that would disappear under proper regularization. To address this concern, we note that when applied to benchmark cases, our method performs well in predicting phase boundaries. Specifically, we have tested it on the Anderson model on the RRG [420] and on the Rosenzweig-Porter model [281]. For the Anderson model on the RRG, the method locates the transition at $W_c \approx 18$, very close to the exact value. We also identify an intermediate region at accessible system sizes where $e^{\mathbb{E}[\ln T_0]}$ decreases exponentially but the system will eventually delocalize due to rare outliers of \mathcal{T}_0 . Importantly, we observe a drift of this intermediate region toward larger disorder values as L increases, consistent with the drift observed in numerical exact diagonalizations. For the Rosenzweig-Porter model, the method accurately identifies all three phases—localized, fractal, and delocalized—even at relatively small system sizes. In summary, while our method might in principle overestimate the effect of large resonances in the intermediate phase, this issue does not appear in the two benchmark cases we have studied.

By studying the model in two different computational bases provides another way to probe the limitations of our method. In particular, we examined how the structure of random initial conditions differs in the spin and Anderson bases. This difference gives rise to significant finite-size effects at intermediate disorder. Nevertheless, the overall physical picture remains robust and consistent with recent findings. Notably, the finite-size phase diagram obtained in the Anderson basis closely matches the results of Refs. [279, 280], where alternative approaches based on real space observables—such as longitudinal spin-spin correlation functions—were employed.

Surprisingly, in our case, finite-size effects manifest differently than in the random-field Ising model in a transverse field previously studied in Ref. [281]. In that model, the method vields a critical disorder strength for the MBL transition that remains approximately stable with increasing system size. In contrast, in the present study, we observe a systematic drift of the critical disorder strength for the MBL transition towards lower values as the system size L increases. Simultaneously, the apparent boundary of the ergodic phase shifts towards higher disorder strengths with growing L. This concomitant trend suggests that the two crossover lines may ultimately merge into a single critical line in the thermodynamic limit, signaling a direct transition from the ergodic phase to the MBL phase. In this scenario, the broad intermediate regime—where delocalization is mediated by rare, system-spanning resonances—would then correspond to a finite-size prethermal crossover that disappears at large L, as proposed in Ref. [257]. The origin of the discrepancy between the finitesize behavior observed here and in Ref. [281] is both intriguing and not yet understood. Further investigation is necessary to clarify the mechanisms behind these differing trends. In particular, it would be highly interesting to investigate whether the finite-size effects reported here can be interpreted within the framework of the renormalization-group flow recently proposed in Refs. [444, 445].

To resolve the spatial structure and statistics of rare resonances, we investigate the progressive delocalization from a random initial state as a function of the distance between the initial and target configurations in Hilbert space. In the strong disorder regime, deep in

the MBL phase, we find that the maximum reachable distance decreases with increasing disorder. As a consequence, at strong disorder, resonant transmission processes remain confined to a small region of Hilbert space around the initial state, allowing only for partial delocalization involving a limited fraction of the degrees of freedom. The extent of this region decreases with increasing system size, reflecting the progressive suppression of rare long-range resonances within the MBL phase. At weak disorder, in the ergodic regime, we observe that uniform spreading of the wavepacket is recovered only at large distances, while at shorter distances, strong resonances responsible for delocalization do not uniformly cover Hilbert space at that scale.

At large disorder, the spatial structure of high-energy eigenstates in Hilbert space exhibits pronounced fluctuations across different disorder realizations, reflecting the difference between typical and rare configurations of the disorder. To probe this heterogeneity, we identified and visualized the dominant resonant paths that form on the Hilbert space graph. This analysis, inspired by approaches to inhomogeneous quantum transport in real space networks [287–289], offers a novel perspective on the MBL instability: it stems from the inclusion of rare resonant paths, which, however, become progressively shorter and increasingly scarce deep within the MBL phase.

This picture of rarefied transmission paths in Hilbert space presents an intriguing direction for further exploration. The original analogy with Anderson localization in two dimensions [289] expands into how these dominant paths pinned by disorder can change suddenly and abruptly producing avalanches—as conceived in the classical setting of directed polymers—when the energy is varied. The depinning transition of the polymers through avalanches can be directly related to the singular behavior of the overlap correlation function between eigenstates at different energies which, in our case, corresponds to the correlation between the Hilbert space Landauer transmissions at different values of the energy—for a given disorder realization. It would be therefore interesting to investigate whether some signatures of these avalanches and shocks are present also in the quantum problem.

Although the importance of system-wide resonances in determining the properties of the MBL transition and the stability of the localized phase has been highlighted in this and several related works [251, 257, 258, 265, 266, 270, 273, 278–281], a proper characterization of the disorder realizations that lead to the formation of these resonances remains an open question. A first step in this direction was the proposal that such many-body resonances manifest as nearly degenerate *cat states* [265, 267, 272], a hypothesis recently tested in Ref. [438]. The anatomy of these nearly degenerate eigenstates reveals resonant events whose probability decreases with increasing disorder strength, consistent with the findings of the present work. Furthermore, rare disorder realizations at intermediate disorder, whose spectra contain such sparse, nearly degenerate eigenstate pairs, also exhibit probabilities to delocalize from a random initial state that are classified as 'rare events' under our metric, appearing in the tails of the distribution of \mathcal{T}_0 .

However, a proper characterization of the structure of rare disorder realizations that give rise to anomalously large delocalization probabilities remains to be performed. This task is highly computationally demanding when using standard sampling techniques. A promising direction would be to employ importance sampling strategies [446]: by biasing the sampling towards disorder configurations that enhance the likelihood of rare resonances, one could develop a genuine large-deviation framework and obtain a more accurate statistical characterization of the spatial structure of these rare events.

Similarly, understanding whether rare delocalization events are favored by the presence of extended regions with anomalously weak disorder—as suggested by the avalanche scenario [100, 243, 248–254]—remains an open problem. To make progress in this direction, it would be valuable to apply our method to study the system's response when coupled to a

thermal bath [100, 251, 257, 258, 260, 447, 448]. Such an analysis could help reveal the signatures in the Hilbert space propagators of rare ergodic bubbles in real space.

The methods proposed here can also be applied to other systems, for example, models of interacting fermions in a quasi-periodic potential, similar to the one realized in cold-atom experiments [361, 449, 450]. In this case, the only source of randomness comes from the choice of the initial state. It would be useful to compare the statistics of rare resonances found in the quasi-periodic case with those of uncorrelated random fields. This could help discriminate between the effects of rare resonances created by large segments with anomalously small values of the disorder in real space, and those due to rare paths with anomalously strong transmission amplitudes in Hilbert space.

A1 Details of the numerical simulations

The task at hand reduces to calculate the 0f elements for the resolvent matrix

$$\mathcal{G}_{0f}(\overline{E}) = (\overline{E}\mathbb{I} - \mathcal{H})_{0f}^{-1}, \qquad (139)$$

for the energy \overline{E} in the middle of the spectrum. We have approximated this average energy by $\overline{E} = \frac{\text{Tr}\mathcal{H}}{\mathcal{N}}$. The initial conditions denoted by '0' are selected by choosing basis states with energy close to \overline{E} . Numerically, we do this by selecting the states $|0\rangle$ following

$$\{|0\rangle\} \equiv \left\{|0\rangle : E_{|0\rangle} = \mathcal{H}_{00} \in \left[\overline{E} - \frac{\overline{E}}{\eta}, \overline{E} + \frac{\overline{E}}{\eta}\right]\right\},$$
 (140)

where η has been chosen to be $\eta=64$ for intermediate and large sizes $L\geq 14$, while $\eta=32$ for the smallest ones, $L\leq 12$. As seen in Eq. (139), we could invert the whole matrix $\overline{E}\mathbb{I}-\mathcal{H}$ and extract the entries of interest i.e. the portion of columns '0' associated to the chosen initial conditions. However, inverting the full matrix of size $\mathcal{N}\times\mathcal{N}$ is computationally expensive. Instead, we directly compute portion of interest by solving the linear system:

$$(\overline{E}\mathbb{I} - \mathcal{H})\mathcal{Y}_{|0\rangle} = \delta_{|0\rangle} , \qquad (141)$$

where $\delta_{|0\rangle}$ is a vector of zeros except for the entry corresponding to the chosen initial condition $|0\rangle$, which is set to one. This linear system is solved using the MUMPS [451, 452] or Pardiso [453] libraries, in their Julia interfaces MUMPS.jl and Pardiso.jl, respectively. We have performed these calculations in both spin and Anderson bases. The number of initial conditions N_0 , as well as the number of samples N_S of the disordered fields h_i is given in the table below. They are presented as a function of the system-size L. The total number of samples, $N_{\text{total}} = N_S \times N_0$, over which the average $\mathbb{E}[\cdots]$ is computed, is also shown—approximately—in the last column. In certain cases—particularly for $\Delta = 1$ in the spin basis—we increased the number of samples where it was deemed necessary.

L	N_0	N_S	$N_{ m total}$
12	$2^{L/2-2}$	125952	2×10^{6}
14	$2^{L/2-2}$	16384	5.25×10^5
16	$2^{L/2-2}$	5120	3.30×10^{5}
18	$2^{L/2-2}$	1280	1.60×10^{5}
20	$2^{L/2-2}$	512	1.30×10^{5}
22	$2^{L/2-1}$	128	1.30×10^5

Table 2: Simulation values in the spin basis

L	N_0	N_S	$N_{ m total}$
8	$2^{L/2-2}$	1638400	6.50×10^{6}
12	$2^{L/2-1}$	65536	2×10^{6}
16	$2^{L/2-1}$	2048	2.60×10^{5}
20	$2^{L/2-1}$	64	3.30×10^{4}

Table 3: Simulation values in the Anderson basis

Moreover, for system sizes that are not divisible by four (e.g., $L=14,\ 18,\ 22$), there are no basis states that exactly satisfy the equator conditions, i.e., $q^{S,A}=0$. Consequently, we

perform an interpolation by averaging \mathcal{T}_0 —and its logarithm, $\ln \mathcal{T}_0$ —between the values computed using the nearest states to the equator. These target states correspond to basis states with overlaps

$$q_{0f}^{S,A} = \pm \frac{2}{L} \,, \tag{142}$$

we select these two sets of target states, and average the quantities of interest between them. This procedure yields a consistent estimate of \mathcal{T}_0 (and $\ln \mathcal{T}_0$) at the equator for system sizes where exact equatorial states are not available.

A2 The computation of the error bars

The error bars shown in Fig. 3.12 are directly extracted from the variances $\mathbb{E}[\mathcal{T}_0^2] - \mathbb{E}[\mathcal{T}_0]^2$ and $\mathbb{E}[(\ln \mathcal{T}_0)^2] - \mathbb{E}[\ln \mathcal{T}_0]^2$, and propagated accordingly for the functional forms of ϕ_a and ϕ_q . For the annealed free-energy, the variance beyond β_\star is ill-defined, making the error bars near β_\star to be numerically unreliable. As a result, the propagated errors associated with the interpolated values of β_\star and $\phi_a(\beta_\star)$ are extremely large and physically meaningless. Therefore, we have employed an alternative approach to assess the reliability of the numerical results produced by our method.

We assess the stability of the relevant quantities— β_{\star} , $\phi_a(\beta_{\star})$, and $\phi_q(\beta=2)$ —under cumulative averaging. In other words, we calculate the relevant quantities with a cubic spline interpolation when averaged over N_S disorder realizations—each of them with N_0 initial conditions—and we keep track of their behavior upon increasing N_S . This is shown in Fig. 23 for both spin (top panels) and Anderson (bottom panels) bases, for the values of disorder width, interaction parameter and sizes shown in the keys.

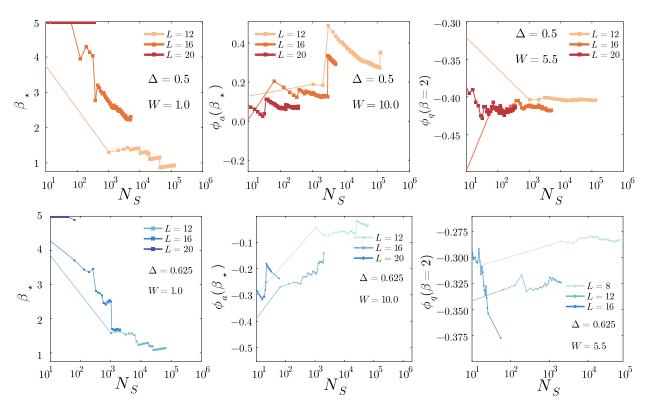


Fig. 23: The values of β_{\star} , $\phi_a(\beta_{\star})$, and $\phi_q(\beta=2)$ as a function of the cumulative number of disorder realizations N_S over which the average $\mathbb{E}[\cdots]$ is taken over. For both spin (top panels) and Anderson (bottom) bases.

We observe that the data exhibit jumps whose size and frequency decrease with increased sampling. The error bars are computed using the last range of values prior to the final average, which we have chosen to be the second half of the cumulant sample sequence.

For example, in the case of L=20 in the Anderson basis, where we use $N_S=64$ disorder realizations, we store the values of $\beta_\star(N_S)$, $\phi_a(\beta_\star,N_S)$, and $\phi_q(\beta=2,N_S)$ corresponding to the cumulative averages for $N_S=1,2,\ldots,N_S=64$. We then consider the second half of this sequence, i.e., from $N_S=32$ to $N_S=64$, and compute the error bars as the difference between the maximum and minimum values of β_\star , $\phi_a(\beta_\star)$, and $\phi_q(\beta=2)$ within this range. Note that β_\star and $\phi_a(\beta_\star)$ are obtained via cubic spline interpolation, for each average over the N_S samples. This measure attempts to assess the stability of the values β_\star , $\phi_a(\beta_\star)$, and $\phi_q(\beta=2)$ upon increasing the sampling of the averages $\mathbb{E}[\mathcal{T}_0(\beta)]$ and $\mathbb{E}[\ln \mathcal{T}_0(\beta)]$.

After obtaining the associated error bars for the values β_{\star} , $\phi_a(\beta_{\star})$, and $\phi_q(\beta=2)$ in this way, we proceed to obtain their respective critical disorder widths $W_{\rm ergo}$, $W_{\rm MBL}$ and $W_{\rm MBL}^{\rm typ}$. The errors for $W_{\rm ergo}$ and $W_{\rm MBL}^{\rm typ}$ are obtained from standard propagation of the errors. We perform a linear interpolation among the values closest to $\beta_{\star}=2$ and $\phi_q(\beta=2)=0$, and propagate the errors accordingly.

On the other hand, when determining $W_{\rm MBL}$, the variation of $\phi_a(\beta_\star,W)$ near the point where $\phi_a(\beta_\star,W)=0$ is very small. This variation is negligible compared to the spacing along the W-axis, which is $\Delta W=1.5$. As a result, the data do not effectively constrain the parameters. This issue is commonly referred to as a *flat direction* in parameter space, or a degeneracy among parameters in non-linear statistical models.

Such degeneracies cause the covariance matrix derived from error propagation to have nearly zero eigenvalues, rendering it highly unstable and non-invertible. To address this, we instead estimate the uncertainty using a bootstrap Monte Carlo resampling approach. In this method, each data point of $\phi_a(\beta_\star,W)$ is randomly perturbed within its error bar, and for each perturbed dataset, the value of W at which $\phi_a(\beta_\star,W)=0$ is recalculated. The standard deviation of the new resulting data values for $W_{\rm MBL}(L)$ provides the error estimate. We use 10^4 resampling iterations in this procedure.

A3 The DPRM correlations

The mapping between the partition functions of directed polymers $\mathcal{Z}_N(\beta)$, defined in Eq. (3.119), and the biased Hilbert space Landauer transmissions $\mathcal{T}_0(\beta)$ is made explicit through the identification of

$$E_{\mathcal{D}} = -\ln|\mathcal{G}_{0f}|\,,\tag{143}$$

where \mathcal{P} therefore corresponds to a 'path' defined between vertices 0 and f in the Hilbert space graph. As discussed in Sec. 3.5.2, the presence of shared edges among different paths \mathcal{P} and \mathcal{P}' introduces correlations between their respective energies $E_{\mathcal{P}}$ and $E_{\mathcal{P}'}$. For the quantum many-body problem the connected correlation $\langle E_{\mathcal{P}}E_{\mathcal{P}'}\rangle_c$ is thus identified with $\langle \ln |\mathcal{G}_{0f}| \ln |\mathcal{G}_{0f'}| \rangle_c$, where the polymers \mathcal{P} and \mathcal{P}' are associated to the Hilbert space 'paths' between vertices 0 to f, and 0 to f', respectively. We measure this connected correlation as a function of the correlation distance—the rescaled overlaps—defined as

$$\zeta_{0f}^{S,A} = 1 - q_{0f}^{S,A} \ . \tag{144}$$

In this case, $|f\rangle$ will vary according to the correlation distance to the initial state $|0\rangle$. The connected correlation is then computed as follows:

$$\langle \ln |\mathcal{G}_{0f}| \ln |\mathcal{G}_{0f'}| \rangle_c \equiv \mathbb{E}[\ln |\mathcal{G}_{0f}| \ln |\mathcal{G}_{0f'}|] - \mathbb{E}[\ln |\mathcal{G}_{0f}|]^2 , \qquad (145)$$

where the average $\mathbb{E}[\cdots]$ is computed for several initial conditions, disorder realizations and among different nodes $|f\rangle$ within the same correlation distance from the chosen initial condition.

In Fig. 24 we show the connected 'energy' correlations $\langle \ln | \mathcal{G}_{0f} | \ln | \mathcal{G}_{0f'} | \rangle_c$ as a function of the correlation distance $\eta_{0f}^{S,A}$. For small disorder (W=1), the correlations remain uniform, appearing as a plateau in both bases. This behavior reflects the fact that, at low disorder, the system is ergodic and there is a proliferation of Hilbert space paths that enable delocalization. At stronger disorders (W=8 and W=16), the connected correlations between 'energies' increase significantly, reflecting the O(1) preferred paths that extend far away in Hilbert space and allow for transmission events.

In the original classical problem of directed polymers in infinite dimensional graphs, this correlation grows linearly with the real space distance of the extending polymer. The exact behaviour in our case is difficult to assess definitively, as the analogous 'polymer' in Hilbert space has a length L/4, that for the largest system size with no interpolation (L=20) corresponds to five flip-flop events in the chain.

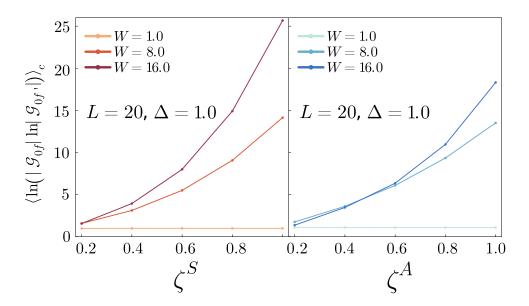


Fig. 24: Correlations of the equivalent polymer energy for the XXZ model. For both spin (left) and Anderson (right) bases, as a function of their respective correlation discante $\zeta^{S,A}$.

A4 Supplementary results for the Anderson basis

In this Section, we present results for the Anderson basis that were omitted from the main text. Specifically, they include: (i) An example of the curves where the characteristic length is extracted for $\beta_{\star}(\zeta^A)=2$ and $\phi_a(\beta_{\star}(\zeta^A))=0$, at the largest size L=20. This is equivalent to Fig. 3.16 for the Anderson basis. And (ii) the eigenstate decays with respect to the most probable basis state, presented in Sec. 3.7,

A4.1 Dependence on the target basis states

Here, we present examples of the dependence of β_{\star} and $\phi_a(\beta_{\star})$ as a function of the correlation distance ζ^A , for L=20. The crossings $\beta_{\star}(\zeta^A)=2$ and $\phi_a(\beta_{\star},\zeta^A)=0$ define the characteristic correlation distances shown on the right panel of Fig. 3.17 for L=20. An equivalent process was performed for L=12,16.

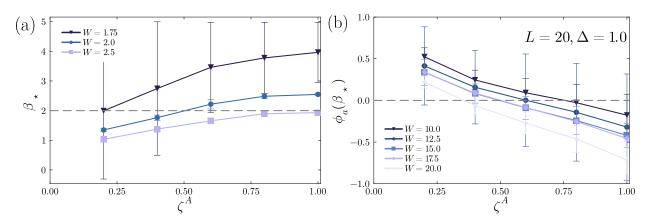


Fig. 25: Calculation of (a) β_{\star} and (b) $\phi_a(\beta_{\star})$ as functions of the correlation distance ζ^S , for $\Delta=1$ and L=20. The values of the disorder widths considered are shown in the legend. Horizontal gray dashed lines indicate the reference values $\beta_{\star}=2$ and $\phi_a(\beta_{\star})=0$. These lines are used to extract the corresponding characteristic correlation distances.

A4.2 The eigenstate amplitude decays

In this case, we did not perform an exact diagonalization. Instead, we extracted the same eigenstates shown in Fig. 3.19 and rotated them using the transformation to the Anderson basis (see Eq. (3.80)).

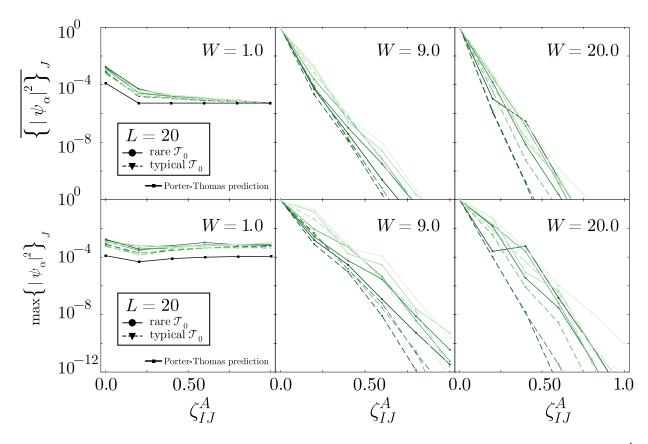


Fig. 26: Decay of the basis state probability as a function of the correlation distance ζ_{IJ}^A between the most probable state I and the other states J that share the same correlation distance ζ^A . Equivalent construction as in Fig. 3.19 but using the Anderson basis.

Conclusions and Perspectives

The central theme of this thesis has been the exploration of numerical and theoretical methods for addressing challenges posed by various ergodicity-breaking mechanisms in physical systems exhibiting quenched disorder.

Originally developed in distinct contexts, these methods have driven important breakthroughs in their respective domains: (i) The mapping between percolation and thermodynamic transitions in FK–CK clusters [50–52] is a well-established framework for characterizing the critical behavior of the ferromagnetic Ising model and its lattice gas counterpart, as well as for designing efficient simulation algorithms to sample critical equilibrium configurations [53–55]; (ii) The SWAP method, first introduced for glass-forming particle systems, has dramatically improved the sampling of equilibrium configurations in low temperature regimes that would otherwise be inaccessible; (iii) Finite-size analyses of quenched and annealed free-energy densities have long been employed to identify the freezing transition β_{\star} in mean-field disordered systems [433, 434].

In this thesis, these approaches have been adapted and applied to new classes of systems with quenched disorder: (i) examining the behavior of FK–CK clusters in the presence of frustration; (ii) extending the SWAP algorithm to spin systems, with emphasis on its impact in frustrated cases; (iii) employing finite-size analyses of 'biased Hilbert space Landauer transmissions' $\mathcal{T}_0(\beta)$ —in both its quenched and annealed versions—in order to account for the influence of rare, long-range many-body resonances in the determination of the critical disorder width $W_{\rm MBL}(L)$ [281].

Applying these tools in new contexts has not only deepened our understanding of the physical systems under study but also offered new insights into the methods themselves. Beyond the methodological transfer, this work underscores the profound physical connections between physical systems that might at first appear unrelated: for example, the parallels between localization transitions and directed polymers [289, 350, 419, 425, 427, 430], and between spin glasses and structural glasses [19, 21, 454]. These connections illustrate the unified phenomenology of disordered systems, where progress in one domain can shed light on challenges in another.

Summary of main results

Chapter 1: The Critical Clusters of Frustrated Spin Systems. We revisited the Fortuin-Kasteleyn-Coniglio-Klein (FK–CK) cluster construction and extended it to the frustrated random-bond Ising model on the Bethe lattice. By analytically continuing the FK–CK bond probabilities to negative values, we demonstrated that these clusters continue to encode the correct thermodynamic critical behavior near the paramagnetic–ferromagnetic transition, even under frustration. Importantly, we identified fundamental limitations for algorithmic implementations: the presence of negative statistical weights renders direct cluster-based Monte Carlo updates infeasible, despite the clusters remaining physically meaningful. We further explored alternative cluster definitions (e.g., α -clusters) and showed that commonly used indicators—such as coinciding percolation and thermodynamic transitions—can be misleading, as they may be artificially reproduced without carrying the genuine critical

correlations associated to the transition. This chapter thus provides both a constructive generalization of FK–CK clusters to frustrated systems and a cautionary perspective on attempts to formulate cluster algorithms in such settings.

Chapter 2: SWAP Dynamics for Frustrated Spin Systems. Inspired by advances in structural glass simulations, we adapted the SWAP algorithm to spin-lattice systems by introducing an auxiliary Δ -model in which Ising spins are assigned continuous length variables. Non-local exchange moves of these lengths accelerate equilibration in two-dimensional Edwards–Anderson spin glasses, particularly at low temperatures, where standard dynamics relaxes extremely slowly. Compared to parallel tempering, our approach achieves a substantial speedup in sampling ground states of the Δ -model, especially when combined with temperature annealing schemes. In contrast, no improvement is observed for non-frustrated ferromagnetic systems, where energy barriers are simpler and equilibration proceeds efficiently via local dynamics alone. However, the acceleration observed here is considerably more modest than that reported in the original particle systems, suggesting that the dynamical constraints present in the latter may indeed play a fundamental role in the much larger speedup observed in structural glasses.

Chapter 3: The Importance of Rare Events in Many-Body Localization. We investigated the many-body localization (MBL) transition in the random-field XXZ chain using a large-deviation framework developed in analogy with mean-field disordered glassy systems [281]. This method isolates the contribution of rare, long-range resonances in Hilbert space by introducing an auxiliary parameter β that tunes the weight of extreme events in the propagator statistics. Applying this framework, we identified three regimes in the (W, Δ) phase diagram: (i) an ergodic phase characterized by uniform delocalization, (ii) an intermediate regime dominated by rare long-range resonances that mediate delocalization in atypical disorder realizations, and (iii) a genuinely localized phase stable against such rare events. The finite-size phase diagram exhibits systematic drifts of the crossover lines with system size, suggesting that in the thermodynamic limit the intermediate regime may collapse into a direct Ergodic–MBL transition. Furthermore, by visualizing resonant transmission paths on the Hilbert space graph, we connected the delocalization mechanism to rare pathways reminiscent of those found in classical directed polymer models [421, 425], reinforcing the analogy between glassy and localization physics [289, 350].

Future directions

The construction of clusters in frustrated systems undergoing a paramagnetic to ferromagnetic transition cannot be implemented numerically using the original definition with the bond probability p_B , since p_B becomes negative. This prevents the application of Wolff- or Swendsen–Wang-type algorithms [53–55], which would otherwise accelerate the dynamics. Nevertheless, the inability to generate physically relevant clusters through standard constructive schemes does not imply that they cannot be generated at all. Alternative approaches remain possible. One promising direction is the use of AI-based generative models or other machine-learning techniques to *learn*, rather than explicitly construct, the clusters. Exactly solvable models on Bethe lattices provide particularly valuable benchmarks for testing such methods. Another benchmark of this approach can be extended into two-dimensional

models. These can be solved exactly using transfer matrix techniques for systems of finite transverse size, allowing for the construction physical clusters, albeit with finite-size effects.

Our results show that critical clusters may only form between parallel spins. In systems with antiferromagnetic couplings, however, configurations where parallel spins are connected by such couplings are strongly suppressed by the Boltzmann weight at low temperatures. As a result, they contribute very little to the spin–spin correlation functions. Paradoxically, these same configurations can still make a significant contribution to the cluster correlation functions. If a bond is present, the corresponding cluster weight may be negative; if the bond is absent, the bonding probability can even exceed one, implying a large effective weight. Such configurations therefore cannot be ignored in the cluster representation. In other words, generating physically meaningful clusters requires sampling rare spin configurations that have extremely low statistical weight.

Although there is no formal proof that any cluster construction in spin glasses reproduces the relevant thermal correlations, the possibility of artificially defining clusters that appear to coincide with the transition remains. While the detection of the spin glass transition temperature T_{SG} —independently of ρ —by tuned multiple-replica clusters is indeed intriguing, it should be interpreted with care, as it may represent yet another artificial smoking gun rather than a genuine physical signature.

In the context of the SWAP method, this encouraging result for two-dimensional systems opens a clear path toward adapting the algorithm to higher-dimensional models. Notable candidates include the 3d Edwards–Anderson (EA) model and the EA model on random regular graphs (RRG), both of which exhibit a spin-glass phase at nonzero temperatures. Investigating the impact of the SWAP method in these cases could yield particularly valuable insights. The RRG case is especially appealing, as the critical temperature of the Δ -model is expected to be well approximated by calculations on the Bethe lattice. This, in turn, would allow for a detailed assessment of the dependence $T_c(\Delta)$ in finite-dimensional systems.

Although the overall magnitude of frustration is reduced, the probability distribution of the converged couplings \mathcal{J}_{ij}^* remains symmetric and does not seem to represent an "easier instance" compared to standard choices such as Gaussian or bimodal distributions. A fuller characterization of these random bonds would require estimating the associated stiffness exponent, known to be $\theta = -0.2793(3)$ for continuous distributions [237, 238]. In the present work, we could not estimate this exponent directly, since the ground states generated by SWAP correspond to different \mathcal{J}_{ij}^* instances, preventing the standard domain-wall analysis of ground-state energies. Nevertheless, it could still be obtained by computing ground states of the \mathcal{J}_{ij}^* instances produced by SWAP using conventional techniques. Pursuing this line of investigation would firmly establish the validity of SWAP in producing non-pathological spin-glass samples.

The influence of rare disorder realizations on the stability of the MBL transition is now well established. However, a proper characterization of the structure of rare disorder realizations that give rise to anomalously large delocalization probabilities remains to be performed. This task is highly computationally demanding when using standard sampling techniques. A promising direction would be to employ importance sampling strategies [446]: by biasing the sampling towards disorder configurations that enhance the likelihood of rare resonances, one could develop a genuine large-deviation framework and obtain a more accurate statistical characterization of the spatial structure of these rare events.

Similarly, understanding whether rare delocalization events are favored by the presence of extended regions with anomalously weak disorder—as suggested by the avalanche scenario [100, 243, 248–254]—remains an open problem. To make progress in this direction, it would be valuable to apply our method to study the system's response when coupled to a thermal bath [100, 251, 257, 258, 260, 447, 448]. Such an analysis could help reveal the

signatures in the Hilbert space propagators of rare ergodic bubbles in real space.

The picture of rarefied transmission paths in Hilbert space offers a compelling lens on the many-body localization (MBL) transition. Building on the analogy with Anderson localization in two dimensions [289], disorder-pinned dominant paths may reorganize abruptly, producing avalanches—akin to those of directed polymers—when the energy is varied. In our framework, such depinning events correspond to singular changes in the correlations between Hilbert-space Landauer transmissions at different energies. Identifying these avalanches and shocks in the quantum many-body problem remains an open challenge.

Finally, the picture of rarefied transmission paths in Hilbert space presents an intriguing direction for further exploration. The original analogy with Anderson localization in two dimensions [289] expands into how these dominant paths pinned by disorder can change suddenly and abruptly producing avalanches—as conceived in the classical setting of directed polymers—when the energy is varied. The depinning transition of the polymers through avalanches can be directly related to the singular behavior of the overlap correlation function between eigenstates at different energies which, in our case, corresponds to the correlation between the Hilbert space Landauer transmissions at different values of the energy—for a given disorder realization. It would be therefore interesting to investigate whether some signatures of these avalanches and shocks are present also in the quantum many-body problem.

Bibliography

- [1] C. C. W. Taylor, *The Atomists: Leucippus and Democritus: Fragments*, Critical edition and commentary on the surviving fragments of the early Greek atomists. (University of Toronto Press, Toronto, 1999).
- [2] A. C. Graham, *Later Mohist Logic, Ethics and Science*, Authoritative analysis of the Mohist Canon, including their proto-scientific ideas about matter and minimal parts. (Chinese University Press, Hong Kong, 1978).
- [3] J. Needham, *Science and Civilisation in China, Volume 2: History of Scientific Thought*, Classic study discussing early Chinese natural philosophy, including Mohist ideas on minimal constituents of matter. (Cambridge University Press, Cambridge, 1956).
- [4] D. Chattopadhyaya, *History of Science and Technology in Ancient India*, Comprehensive overview of Indian scientific thought, including Vaisheshika atomic theory. (Firma KLM, Calcutta, 1991).
- [5] S. Radhakrishnan, *Indian Philosophy, Volume 2*, Detailed discussion of Vaisheshika and Nyaya schools and their atomic theories. (Oxford University Press, Oxford, 1951).
- [6] H. Bhattacharyya, *The Cultural Heritage of India, Volume 1: Philosophy* (The Ramakrishna Mission Institute of Culture, 1978).
- [7] K. C. Shastri, *Jaina Philosophy: Historical and Comparative Study*, Covers Jaina metaphysics and atomic theory. (Bhogilal Leherchand Institute, 1964).
- [8] S. H. Nasr, Science and Civilization in Islam, Overview of Islamic scientific thought, including Ash'arite discussions of atomism. (Harvard University Press, Cambridge, 2007).
- [9] H. A. Wolfson, "The Kalam Argument for Creation in Medieval Jewish Philosophy", The Harvard Theological Review **66**, Examines Ash'arite and Kalam atomistic theories and their influence., 153–173 (1976).
- [10] S. G. Brush, Statistical Physics and the Atomic Theory of Matter: From Boyle and Newton to Landau and Onsager (Princeton University Press, 1983).
- [11] M. J. Klein, *Thermodynamics in Einstein's Universe: Statistical Mechanics and the Foundations of Physics* (Princeton University Press, 1986).
- [12] P. W. Anderson, "More Is Different: Broken Symmetry and the Nature of the Hierarchical Structure of Science", Science 177, 393–396 (1972).
- [13] C. Cercignani, *Ludwig Boltzmann: The Man Who Trusted Atoms*, A biographical and historical account of Ludwig Boltzmann and the development of statistical mechanics. (Oxford University Press, Oxford, 1998).
- [14] V. I. Arnold and A. Avez, *Ergodic Problems of Classical Mechanics* (W. A. Benjamin, New York, 1968).
- [15] D. V. Anosov, *Geodesic flows on closed Riemannian manifolds with negative curvature*, Vol. 90 (Proceedings of the Steklov Institute of Mathematics, 1967).

- [16] N. Simányi, "Proof of the Boltzmann-Sinai Ergodic Hypothesis for typical hard ball systems", Inventiones Mathematicae **154**, 123–178 (2003).
- [17] N. Simányi, "Proof of the Boltzmann–Sinai Ergodic Hypothesis for typical hard ball systems on a torus", Annals of Mathematics **161**, 2037–2066 (2013).
- [18] Y. G. Sinai, "On the foundation of the ergodic hypothesis for a dynamical system of statistical mechanics", Soviet Mathematics Doklady 4, 1818–1822 (1963).
- [19] L. Berthier and G. Biroli, "Theoretical perspective on the glass transition and amorphous materials", Rev. Mod. Phys. **83**, 587–645 (2011).
- [20] L. Berthier, G. Biroli, J.-P. Bouchaud, and G. Tarjus, "Can the glass transition be explained without a growing static length scale?", The Journal of Chemical Physics 150 (2019).
- [21] A. Cavagna, "Supercooled liquids for pedestrians", Physics Reports 476 (2009).
- [22] J. A. Mydosh, Spin glasses: an experimental introduction (CRC press, 1993).
- [23] M. A. Ruderman and C. Kittel, "Indirect exchange coupling of nuclear magnetic moments by conduction electrons", Physical Review **96**, 99–102 (1954).
- [24] T. Kasuya, "A theory of metallic ferro- and antiferromagnetism on Zener's model", Progress of Theoretical Physics **16**, 45–57 (1956).
- [25] K. Yosida, "Magnetic properties of Cu-Mn alloys", Physical Review **106**, 893–898 (1957).
- [26] S. F. Edwards and P. W. Anderson, "Theory of spin glasses", Journal of Physics F: Metal Physics **5**, 965–974 (1975).
- [27] D. Sherrington and S. Kirkpatrick, "Solvable Model of a Spin-Glass", Phys. Rev. Lett. **35**, 1792–1796 (1975).
- [28] S. Kirkpatrick and D. Sherrington, "Infinite-ranged models of spin-glasses", Phys. Rev. B 17, 4384–4403 (1978).
- [29] G. Parisi, "Toward a mean field theory for spin glasses", Physics Letters A **73**, 203–205 (1979).
- [30] G. Parisi, "Infinite number of order parameters for spin-glasses", Phys. Rev. Lett. **43**, 1754–1756 (1979).
- [31] G. Parisi, "A sequence of approximated solutions to the S-K model for spin glasses", Journal of Physics A: Mathematical and General **13**, L115–L117 (1980).
- [32] G. Parisi, "The order parameter for spin glasses: a function on the interval 0–1", Journal of Physics A: Mathematical and General **13**, 1101–1112 (1980).
- [33] G. Parisi, "Magnetic properties of spin glasses in a new mean field theory", Journal of Physics A: Mathematical and General **13**, 1887–1895 (1980).
- [34] G. Parisi, "Order parameter for spin-glasses", Phys. Rev. Lett. **50**, 1946–1948 (1983).
- [35] M. Mézard, G. Parisi, and M. A. Virasoro, *Spin Glass Theory and Beyond: An Introduction to the Replica Method and Its Applications* (World Scientific, 1987).
- [36] A. Altieri and M. Baity-Jesi, "An introduction to the theory of spin glasses", in *Encyclopedia of condensed matter physics* (Elsevier, 2024), pp. 361–370.
- [37] F. Guerra and F. L. Toninelli, Communications in Mathematical Physics **230**, 71–79 (2002).

- [38] M. Talagrand, "Replica symmetry breaking and exponential inequalities for the Sherrington-Kirkpatrick model", The Annals of Probability **28** (2000).
- [39] R. Rammal, G. Toulouse, and M. A. Virasoro, "Ultrametricity for Physicists", Rev. Mod. Phys. **58**, 765–788 (1986).
- [40] D. J. Thouless, P. W. Anderson, and R. G. Palmer, "Solution of "Solvable model of a spin glass", Philosophical Magazine **35**, 593–601 (1977).
- [41] L. F. Cugliandolo and J. Kurchan, "Analytical solution of the off-equilibrium dynamics of a long-range spin-glass model", Phys. Rev. Lett. **71**, 173–176 (1993).
- [42] L. F. Cugliandolo and J. Kurchan, "On the out-of-equilibrium relaxation of the Sherrington-Kirkpatrick model", Journal of Physics A: Mathematical and General **27**, 5749–5772 (1994).
- [43] L. F. Cugliandolo and D. S. Dean, "Full dynamical solution for a spherical spin-glass model", Journal of Physics A: Mathematical and General **28**, 4213–4234 (1995).
- [44] L. F. Cugliandolo, "Recent applications of dynamical mean-field methods", Annual Review of Condensed Matter Physics **15** (2023).
- [45] A. J. Bray and S. A. Roberts, "Renormalisation-group approach to the spin glass transition in finite magnetic fields", Journal of Physics C: Solid State Physics 13, 5405–5412 (1980).
- [46] E. D. Dahlberg, I. González-Adalid Pemartín, E. Marinari, V. Martin-Mayor, J. Moreno-Gordo, R. L. Orbach, I. Paga, G. Parisi, F. Ricci-Tersenghi, J. J. Ruiz-Lorenzo, and D. Yllanes, "Spin-glass dynamics: experiment, theory and simulation", arXiv preprint arXiv:2412.08381 (2024).
- [47] D. S. Fisher and D. A. Huse, "Ordered Phase of Short-Range Ising Spin-Glasses", Phys. Rev. Lett. **56**, 1601–1604 (1986).
- [48] S. Puri, ed., *Kinetics of phase transitions*, en, Chemistry (CRS Press, Boca Raton, Fla., 2009).
- [49] R. Bausch, H. K. Janssen, and H. Wagner, "Renormalized field theory of critical dynamics", Zeitschrift für Physik B Condensed Matter and Quanta **24**, 113–127 (1976).
- [50] C. M. Fortuin and P. W. Kasteleyn, "On the random-cluster model: I. Introduction and relation to other models", Physica **57**, 536–564 (1972).
- [51] P. W. Kasteleyn and C. M. Fortuin, "Phase transitions in lattice systems with random local properties", J. Phys. Soc. Jpn. Suppl. **26**, 11 (1969).
- [52] A. Coniglio and W. Klein, "Clusters and Ising critical droplets: a renormalisation group approach", J. Phys. A **13**, 2775 (1980).
- [53] R. H. Swendsen and J.-S. Wang, "Replica Monte Carlo simulation of spin-glasses", Phys. Rev. Lett. **57**, 2607 (1986).
- [54] R. H. Swendsen and J.-S. Wang, "Nonuniversal critical dynamics in Monte Carlo simulations", Phys. Rev. Lett. **58**, 86 (1987).
- [55] U. Wolff, "Collective Monte Carlo updating for spin systems", Phys. Rev. Lett. **62**, 361 (1989).
- [56] H. Li, J. He, and R. L. Orbach, "Nature of temperature chaos in spin glasses", Phys. Rev. B **111** (2025).

- [57] T. Rizzo and A. Crisanti, "Chaos in Temperature in the Sherrington-Kirkpatrick Model", Phys. Rev. Lett. **90** (2003).
- [58] H. Nishimori, M. Ohzeki, and M. Okuyama, "Temperature chaos as a logical consequence of reentrant transition in spin glasses", arXiv preprint arXiv:2507.00276 (2025).
- [59] "The Mpemba effect in spin glasses is a persistent memory effect", Proc. Natl. Acad. Sci. U.S.A. **116**, 15350–15355 (2019).
- [60] H. Rieger, B. Steckemetz, and M. Schreckenberg, "Aging and Domain Growth in the Two-Dimensional Ising Spin Glass Model", en, Europhys. Lett. **27**, 485–490 (1994).
- [61] L. Berthier, D. Coslovich, A. Ninarello, and M. Ozawa, "Equilibrium Sampling of Hard Spheres up to the Jamming Density and Beyond", Phys. Rev. Lett. **116**, 238002 (2016).
- [62] A. Ninarello, L. Berthier, and D. Coslovich, "Models and Algorithms for the Next Generation of Glass Transition Studies", Phys. Rev. X 7, 021039 (2017).
- [63] L. Berthier, E. Flenner, C. J. Fullerton, C. Scalliet, and M. Singh, "Efficient swap algorithms for molecular dynamics simulations of equilibrium supercooled liquids", J. Stat. Mech. **2019**, 064004 (2019).
- [64] T. R. Kirkpatrick and P. G. Wolynes, "Stable and metastable states in mean-field Potts and structural glasses", Phys. Rev. B **36** (1987).
- [65] T. R. Kirkpatrick and D. Thirumalai, "Dynamics of the Structural Glass Transition and the *p*-Spin—Interaction Spin-Glass Model", Phys. Rev. Lett. **58** (1987).
- [66] T. R. Kirkpatrick and P. G. Wolynes, "Connections between some kinetic and equilibrium theories of the glass transition", Phys. Rev. A **35** (1987).
- [67] T. R. Kirkpatrick, D. Thirumalai, and P. G. Wolynes, "Scaling concepts for the dynamics of viscous liquids near an ideal glassy state", Phys. Rev. A **40** (1989).
- [68] T. R. Kirkpatrick and D. Thirumalai, "Random solutions from a regular density functional Hamiltonian: a static and dynamical theory for the structural glass transition", Journal of Physics A: Mathematical and General **22** (1989).
- [69] T. Kirkpatrick and D. Thirumalai, "Colloquium: Random first order transition theory concepts in biology and physics", Rev. Mod. Phys. **87** (2015).
- [70] B. Guiselin, G. Tarjus, and L. Berthier, "Is glass a state of matter?", Phys. Chem. Glasses **63**, 136–144 (2022).
- [71] J. P. Garrahan and D. Chandler, "Coarse-grained microscopic model of glass formers", Proc. Natl. Acad. Sci. U.S.A. **100**, 9710–9714 (2003).
- [72] D. Chandler and J. P. Garrahan, "Dynamics on the Way to Forming Glass: Bubbles in Space-Time", Annual Review of Physical Chemistry **61**, 191–217 (2010).
- [73] F. Ritort and P. Sollich, "Glassy dynamics of kinetically constrained models", Advances in Physics **52** (2003).
- [74] G. H. Fredrickson and H. C. Andersen, "Kinetic Ising Model of the Glass Transition", Phys. Rev. Lett. **53**, 1244–1247 (1984).
- [75] J. Jackle and S. Eisinger, "A hierarchically constrained kinetic Ising model", Zeitschrift für Physik B Condensed Matter **84**, 115–124 (1991).
- [76] M. Wyart and M. E. Cates, "Does a Growing Static Length Scale Control the Glass Transition?", Phys. Rev. Lett. **119**, 195501 (2017).

- [77] R. Gutiérrez, J. P. Garrahan, and R. L. Jack, "Accelerated relaxation and suppressed dynamic heterogeneity in a kinetically constrained (East) model with swaps", Journal of Statistical Mechanics: Theory and Experiment **2019**, 094006 (2019).
- [78] J. M. Deutsch, "Quantum statistical mechanics in a closed system", Phys. Rev. A 43, 2046–2049 (1991).
- [79] M. Srednicki, "Chaos and quantum thermalization", Phys. Rev. E **50**, 888–901 (1994).
- [80] M. Srednicki, "Thermal fluctuations in quantized chaotic systems", Journal of Physics A: Mathematical and General **29**, L75–L79 (1996).
- [81] M. Srednicki, "The approach to thermal equilibrium in quantized chaotic systems", Journal of Physics A: Mathematical and General **32**, 1163–1175 (1999).
- [82] J. M. Deutsch, "Eigenstate thermalization hypothesis", Reports on Progress in Physics **81**, 082001 (2018).
- [83] 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**, 239–362 (2016).
- [84] A. I. Shnirelman, "Ergodic Properties of Eigenfunctions", Uspekhi Matematicheskikh Nauk (Russian Math. Surveys) **29**, 181–182 (1974).
- [85] L. Foini and J. Kurchan, "Eigenstate Thermalization and Rotational Invariance in Ergodic Quantum Systems", Phys. Rev. Lett. **123** (2019).
- [86] S. Pappalardi, L. Foini, and J. Kurchan, "Eigenstate Thermalization Hypothesis and Free Probability", Phys. Rev. Lett. **129** (2022).
- [87] L. Foini and J. Kurchan, "Eigenstate thermalization hypothesis and out of time order correlators", Phys. Rev. E **99** (2019).
- [88] P. W. Anderson, "Absence of Diffusion in Certain Random Lattices", en, Phys. Rev. **109** (1958).
- [89] D. Basko, I. Aleiner, and B. Altshuler, "Metal–insulator transition in a weakly interacting many-electron system with localized single-particle states", Ann. Phys. (N. Y.) **321**, 1126–1205 (2006).
- [90] I. V. Gornyi, A. D. Mirlin, and D. G. Polyakov, "Interacting electrons in disordered wires: Anderson localization and low-temperature transport", Phys. Rev. Lett. **95** (2005).
- [91] D. A. Abanin, E. Altman, I. Bloch, and M. Serbyn, "*Colloquium*: Many-body localization, thermalization, and entanglement", en, Rev. Mod. Phys. **91** (2019).
- [92] F. Alet and N. Laflorencie, "Many-body localization: An introduction and selected topics", en, C. R. Phys. **19** (2018).
- [93] R. Nandkishore and D. A. Huse, "Many-Body Localization and Thermalization in Quantum Statistical Mechanics", Annu. Rev. Condens. Mat. Phys. **6**, 15–38 (2015).
- [94] P. Sierant, M. Lewenstein, A. Scardicchio, L. Vidmar, and J. Zakrzewski, "Many-body localization in the age of classical computing*", Rep. Prog. Phys. **88** (2025).
- [95] J. Šuntajs, J. Bonča, T. Prosen, and L. Vidmar, "Quantum chaos challenges many-body localization", en, Phys. Rev. E **102** (2020).
- [96] J. Šuntajs, J. Bonča, T. Prosen, and L. Vidmar, "Ergodicity breaking transition in finite disordered spin chains", en, Phys. Rev. B **102** (2020).

- [97] M. Kiefer-Emmanouilidis, R. Unanyan, M. Fleischhauer, and J. Sirker, "Evidence for Unbounded Growth of the Number Entropy in Many-Body Localized Phases", en, Phys. Rev. Lett. **124** (2020).
- [98] P. Sierant, M. Lewenstein, and J. Zakrzewski, "Polynomially Filtered Exact Diagonalization Approach to Many-Body Localization", Phys. Rev. Lett. **125** (2020).
- [99] D. Sels and A. Polkovnikov, "Dynamical obstruction to localization in a disordered spin chain", en, Phys. Rev. E **104** (2021).
- [100] D. Sels, "Bath-induced delocalization in interacting disordered spin chains", en, Phys. Rev. B **106** (2022).
- [101] M. Kiefer-Emmanouilidis, R. Unanyan, M. Fleischhauer, and J. Sirker, "Slow delocalization of particles in many-body localized phases", en, Phys. Rev. B **103** (2021).
- [102] P. Sierant and J. Zakrzewski, "Challenges to observation of many-body localization", en, Phys. Rev. B **105** (2022).
- [103] D. Sels and A. Polkovnikov, "Thermalization of Dilute Impurities in One-Dimensional Spin Chains", en, Phys. Rev. X **13** (2023).
- [104] N. Goldenfeld, *Lectures on Phase Transitions and the Renormalization Group* (CRC Press, Mar. 2018).
- [105] V. S. Dotsenko, P. Windey, G. Harris, E. Marinari, E. Martinec, and M. Picco, "Critical and topological properties of cluster boundaries in the 3D Ising model", Phys. Rev. Lett. **71**, 811–814 (1993).
- [106] K. Binder, "Coniglio–Klein droplets: a fruitful concept to understand phase transitions geometrically", Journal of Physics A: Mathematical and Theoretical **50**, 061001 (2017).
- [107] G. R. Grimmett, *The Random-Cluster Model* (Springer, 2006).
- [108] S. Esposito and A. Amabile, "Percolating lives: statistical mechanics in Naples", Eur. Phys. J. H **49**, 12 (2024).
- [109] R. G. Edwards and A. D. Sokal, "Generalization of the Fortuin–Fasteleyn–Swendsen–Wang representation and Monte Carlo algorithm", Phys. Rev. D **38**, 2009 (1988).
- [110] A. Coniglio and A. Fierro, "Correlated Percolation", in *Encyclopedia of complexity and systems science*, edited by R. A. Meyers (Springer, 2009), pp. 1596–1615.
- [111] C. Amitrano, F. di Liberto, R. Figari, and F. Peruggi, "*A-B* droplets for a two-dimensional antiferromagnetic Ising model in external field *H*", J. Phys. A **16**, 3925 (1983).
- [112] N. Jan, A. Coniglio, and D. Stauffer, "Study of droplets for correlated site-bond percolation in two dimensions", J. Phys. A **15**, L699 (1982).
- [113] J.-S. Wang and R. H. Swendsen, "Cluster Monte Carlo algorithms", Physica A **167**, 565–579 (1990).
- [114] M. E. J. Newman and G. T. Barkema, *Monte Carlo Methods in Statistical Physics* (Clarendon Press, 1999).
- [115] D. Landau and K. Binder, *A Guide to Monte Carlo Simulations in Statistical Physics*, 5th ed. (Cambridge University Press, 2021).
- [116] A. Coniglio, F. Di Liberto, G. Monroy, and F. Peruggi, "Cluster approach to spin glasses and the frustrated-percolation problem", Phys. Rev. B **44**, 12605 (1991).

- [117] L. Chayes, J. Machta, and O. Redner, "Graphical Representations for Ising Systems in External Fields", J. Stat. Phys. **93**, 17–32 (1998).
- [118] J. Houdayer, "A cluster Monte Carlo algorithm for 2-dimensional spin glasses", Eur. Phys. J. B **22**, 479–484 (2001).
- [119] J. Machta, C. M. Newman, and D. L. Stein, "The Percolation Signature of the Spin Glass Transition", Journal of Statistical Physics **130**, 113–128 (2007).
- [120] L. Münster and M. Weigel, "Cluster percolation in the two-dimensional Ising spin glass", Phys. Rev. E **107** (2023).
- [121] L. Münster and M. Weigel, "Spin glasses and percolation", Frontiers in Physics **12** (2024).
- [122] H. Fajen, A. K. Hartmann, and A. P. Young, "Percolation of Fortuin–Kasteleyn clusters for the random-bond Ising model", Phys. Rev. E **102**, 012131 (2020).
- [123] P. Charbonneau and M. Tarzia, "Solution of disordered microphases in the Bethe approximation", J. Chem. Phys. **155**, 024501 (2021).
- [124] D. Stauffer and A. Aharony, *Introduction to percolation theory*, en, 2nd ed. (Taylor & Francis, London, England, July 1994).
- [125] A. Aharony and D. Stauffer, *Introduction To Percolation Theory* (Taylor & Francis, 2003).
- [126] M. Mézard and G. Parisi, "The Bethe lattice spin glass revisited", Eur. Phys. J. B **20**, 217–233 (2001).
- [127] M. Mézard and A. Montanari, *Information, Physics, and Computation* (Oxford University Press, 2009).
- [128] M. Ostilli, "Cayley Trees and Bethe Lattices: A concise analysis for mathematicians and physicists", Physica A: Statistical Mechanics and its Applications **391**, 3417–3423 (2012).
- [129] K. Binder, ""Clusters" in the Ising model, metastable states and essential singularity", Annals of Physics **98**, 390–417 (1976).
- [130] H. Müller-Krumbhaar, "Percolation in a lattice system with particle interaction", Physics Letters A **50**, 27–28 (1974).
- [131] A. Coniglio, "Some cluster-size and percolation problems for interacting spins", Phys. Rev. B **13**, 2194–2207 (1976).
- [132] R. H. Swendsen, J.-S. Wang, and A. M. Ferrenberg, "New Monte Carlo methods for improved efficiency of computer simulations in statistical mechanics", The Monte Carlo method in condensed matter physics, 75–91 (1992).
- [133] V. Cataudella, G. Franzese, M. Nicodemi, A. Scala, and A. Coniglio, "Critical clusters and efficient dynamics for frustrated spin models", Phys. Rev. Lett. **72**, 1541 (1994).
- [134] P. W. Leung and C. L. Henley, "Percolation properties of the Wolff clusters in planar triangular spin models", Phys. Rev. B **43**, 752–759 (1991).
- [135] L. de Arcangelis, A. Coniglio, and F. Peruggi, "Percolation transition in spin glasses", Europhys. Lett. **14**, 515 (1991).
- [136] V. Cataudella, "Percolation transition in systems with frustation", Physica A **183**, 249–254 (1992).
- [137] A. Kalz, A. Honecker, S. Fuchs, and T. Pruschke, "Phase diagram of the Ising square lattice with competing interactions", Eur. Phys. J. B **65**, 533 (2008).

- [138] Z. Zhu, A. J. Ochoa, and H. G. Katzgraber, "Efficient Cluster Algorithm for Spin Glasses in Any Space Dimension", Phys. Rev. Lett. **115**, 077201 (2015).
- [139] M. Zheng, M. Tarzia, and P. Charbonneau, "Weakening the critical dynamical slow-down of models with SALR interactions", J. Chem. Phys. **157**, 181103 (2022).
- [140] S. Cho and M. P. A. Fisher, "Criticality in the two-dimensional random-bond Ising model", Phys. Rev. B **55**, 1025–1031 (1997).
- [141] F. Merz and J. T. Chalker, "Two-dimensional random-bond Ising model, free fermions, and the network model", Phys. Rev. B **65**, 054425 (2002).
- [142] A. Honecker, M. Picco, and P. Pujol, "Universality Class of the Nishimori Point in the $2D \pm J$ Random-Bond Ising Model", Phys. Rev. Lett. **87**, 047201 (2001).
- [143] M. Picco, A. Honecker, and P. Pujol, "Strong disorder fixed points in the two-dimensional random-bond Ising model", J. Stat. Mech. **2006**, P09006 (2006).
- [144] M. Hasenbusch, F. Parisen Toldin, A. Pelissetto, and E. Vicari, "Multicritical Nishimori point in the phase diagram of the $\pm J$ Ising model on a square lattice", Phys. Rev. E 77, 051115 (2008).
- [145] M. Hasenbusch, F. Parisen Toldin, A. Pelissetto, and E. Vicari, "Universal dependence on disorder of two-dimensional randomly diluted and random-bond $\pm J$ Ising models", Phys. Rev. E **78**, 011110 (2008).
- [146] F. Parisen Toldin, A. Pelissetto, and E. Vicari, "Strong-disorder paramagnetic ferromagnetic fixed point in the square-lattice $\pm J$ Ising model", J. Stat. Phys. **135**, 1039 (2009).
- [147] Y. Liu, D. Wang, X. Wang, D.-X. Yao, and L.-H. Tang, Magnetization-resolved density of states and quasi-first order transition in the two-dimensional random bond Ising model: an entropic sampling study, arXiv:2505.04298, 2025.
- [148] J. T. Chayes, L. Chayes, J. P. Sethna, and D. J. Thouless, "A mean field spin glass with short-range interactions", Communications in Mathematical Physics **106**, 41–89 (1986).
- [149] J. M. Carlson, J. T. Chayes, L. Chayes, J. P. Sethna, and D. J. Thouless, "Critical Behavior of the Bethe Lattice Spin Glass", EPL (Europhysics Letters) **6**, 475–475 (1988).
- [150] J. M. Carlson, J. T. Chayes, L. Chayes, J. P. Sethna, and D. J. Thouless, "Bethe lattice spin glass: The effects of a ferromagnetic bias and external fields. I. Bifurcation analysis", Journal of Statistical Physics **61**, 987–1067 (1990).
- [151] J. M. Carlson, J. T. Chayes, J. P. Sethna, and D. J. Thouless, "Bethe lattice spin glass: The effects of a ferromagnetic bias and external fields. II. Magnetized spin-glass phase and the de Almeida-Thouless line", Journal of Statistical Physics **61**, 1069–1084 (1990).
- [152] H. Nishimori, "Internal Energy, Specific Heat and Correlation Function of the Bond-Random Ising Model", Prog. Theor. Phys. **66**, 1169–1181 (1981).
- [153] H. Nishimori, *Statistical Physics of Spin Glasses and Information Processing* (Oxford University Press, July 2001).
- [154] R. Agrawal, L. F. Cugliandolo, L. Faoro, L. B. Ioffe, and M. Picco, "Nonequilibrium critical dynamics of the two-dimensional $\pm J$ Ising model", Phys. Rev. E **108**, 064131 (2023).

- [155] R. Agrawal, L. F. Cugliandolo, L. Faoro, L. B. Ioffe, and M. Picco, "Dynamical critical behavior on the Nishimori point of frustrated Ising models", Phys. Rev. E **110**, 034120 (2024).
- [156] V. Dotsenko, *An Introduction to the Theory of Spin Glasses and Neural Networks* (World Scientific, 1994).
- [157] A. Pagnani, G. Parisi, and M. Ratiéville, "Near-optimal configurations in mean-field disordered systems", Phys. Rev. E **68**, 046706 (2003).
- [158] M. Mézard and G. Parisi, Journal of Statistical Physics 111, 1–34 (2003).
- [159] C. Kwon and D. J. Thouless, "Ising spin glass at zero temperature on the Bethe lattice", Phys. Rev. B **37**, 7649–7654 (1988).
- [160] J. Kertész and D. E. Wolf, "Anomalous roughening in growth processes", Phys. Rev. Lett. **62**, 2571–2574 (1989).
- [161] A. Coniglio, F. de Liberto, G. Monroy, and F. Peruggi, "Exact relations between droplets and thermal fluctuations in external field", J. Phys. A: Math. Gen. **22**, L837 (1989).
- [162] J. R. L. d. Almeida and D. J. Thouless, "Stability of the Sherrington-Kirkpatrick solution of a spin glass model", Journal of Physics A: Mathematical and General 11, 983–990 (1978).
- [163] Y. Matsuda, M. Müller, H. Nishimori, T. Obuchi, and A. Scardicchio, "Distribution of partition function zeros of the ±Jmodel on the Bethe lattice", Journal of Physics A: Mathematical and Theoretical **43**, 285002 (2010).
- [164] T. Jörg, "Cluster Monte Carlo Algorithms for Diluted Spin Glasses", en, Progress of Theoretical Physics Supplement **157**, 349–352 (2005).
- [165] T. Jörg, "Critical behavior of the three-dimensional bond-diluted Ising spin glass: Finite-size scaling functions and universality", en, Phys. Rev. B **73**, 224431 (2006).
- [166] C. M. Newman and D. L. Stein, "Short-Range Spin Glasses: Results and Speculations", in *Spin glasses* (Springer Berlin Heidelberg, 2007), pp. 159–175.
- [167] D. J. Earl and M. W. Deem, "Parallel Tempering: Theory, Applications, and New Perspectives", Phys. Chem. Chem. Phys. **7**, 3910 (2005).
- [168] E. Marinari and G. Parisi, "Simulated Tempering: A New Monte Carlo Scheme", Europhys. Lett. **19**, 451 (1992).
- [169] W. Wang, J. Machta, and H. G. Katzgraber, "Population annealing: Theory and application in spin glasses", Phys. Rev. E **92** (2015).
- [170] A. Barzegar, C. Pattison, W. Wang, and H. G. Katzgraber, "Optimization of population annealing Monte Carlo for large-scale spin-glass simulations", Phys. Rev. E **98** (2018).
- [171] M. Weigel, L. Barash, L. Shchur, and W. Janke, "Understanding population annealing Monte Carlo simulations", Phys. Rev. E **103** (2021).
- [172] C. Baldassi, C. Borgs, J. T. Chayes, A. Ingrosso, C. Lucibello, L. Saglietti, and R. Zecchina, "Unreasonable effectiveness of learning neural networks: From accessible states and robust ensembles to basic algorithmic schemes", Proc. Nat. Acad. Sc. 113, E7655 (2016).
- [173] M. C. Angelini and F. Ricci-Tersenghi, "Limits and performances of algorithms based on simulated annealing in solving sparse hard inference problems", Phys. Rev. X 13, 021011 (2023).

- [174] F. Belletti et al., "JANUS: an FPGA-based System for High Performance Scientific Computing", Computing in Science & Engineering 11, 48 (2009).
- [175] M. Baity-Jesi, A. Cruz, L. A. Fernandez, J. Gil-Narvion, A. Gordillo-Guerrero, D. Iñiguez, A. Maiorano, V. Martin-Mayor, J. Monforte-Garcia, A. Muñoz-Sudupe, D. Navarro, G. Parisi, F. Ricci-Tersenghi, J. J. Ruiz-Lorenzo, S. F. Schifano, B. Seoane, A. Tarancon, R. Tripiccione, and D. Yllanes, "Critical behavior of the three-dimensional Ising spin glass in a magnetic field", Phys. Rev. B 89, 144202 (2014).
- [176] C. Fan, M. Shen, Z. Nussinov, Z. Liu, Y. Sun, and Y.-Y. Liu, "Searching for spin glass ground states through deep reinforcement learning", Nat. Commun. 14, 725 (2023).
- [177] S. Boettcher, "Deep reinforced learning heuristic tested on spin-glass ground states: The larger picture", CoRR **abs/2302.10848** (2023).
- [178] L. Berthier and D. R. Reichman, "Modern computational studies of the glass transition", Nat. Rev. Phys. 5, 102 (2023).
- [179] C. Scalliet, B. Guiselin, and L. Berthier, "Thirty milliseconds in the life of a supercooled liquid", Phys. Rev. X **12**, 041028 (2022).
- [180] S. Albert, T. Bauer, M. Michl, G. Biroli, J.-P. Bouchaud, A. Loidl, P. Lunkenheimer, R. Tourbot, C. Wiertel-Gasquet, and F. Ladieu, "Fifth-order susceptibility unveils growth of thermodynamic amorphous order in glass-formers", Science **352** (2016).
- [181] S. Yaida, L. Berthier, P. Charbonneau, and G. Tarjus, "Point-to-set lengths, local structure, and glassiness", Phys. Rev. E **94** (2016).
- [182] L. Berthier and W. Kob, "Static point-to-set correlations in glass-forming liquids", Phys. Rev. E **85** (2012).
- [183] L. Berthier, G. Biroli, J.-P. Bouchaud, and G. Tarjus, "Can the glass transition be explained without a growing static length scale?", J. Chem. Phys. **150**, 094501 (2019).
- [184] D. Chandler and J. P. Garrahan, "Dynamics on the way to forming glass: Bubbles in space-time", Annual review of physical chemistry **61**, 191–217 (2010).
- [185] S. Guchhait and R. Orbach, Phys. Rev. Lett. 112, 126401 (2014).
- [186] A. K. Hartmann, "Ground States of Two-Dimensional Ising Spin Glasses: Fast Algorithms, Recent Developments and a Ferromagnet-Spin Glass Mixture", en, J. Stat. Phys. **144**, 519–540 (2011).
- [187] H. Khoshbakht and M. Weigel, "Domain-wall excitations in the two-dimensional Ising spin glass", Phys. Rev. B **97**, 064410 (2018).
- [188] W. Kauzmann, "The Nature of the Glassy State and the Behavior of Liquids at Low Temperatures.", Chem. Rev. **43**, 219–256 (1948).
- [189] J.-P. Bouchaud and G. Biroli, "On the Adam-Gibbs-Kirkpatrick-Thirumalai-Wolynes scenario for the viscosity increase in glasses", The Journal of Chemical Physics **121**, 7347–7354 (2004).
- [190] A. Montanari and G. Semerjian, "Rigorous Inequalities Between Length and Time Scales in Glassy Systems", Journal of Statistical Physics **125**, 23–54 (2006).
- [191] C. Crauste-Thibierge, C. Brun, F. Ladieu, D. L'Hôte, G. Biroli, and J.-P. Bouchaud, "Evidence of Growing Spatial Correlations at the Glass Transition from Nonlinear Response Experiments", Phys. Rev. Lett. **104** (2010).
- [192] C. K. Thomas and A. A. Middleton, "Numerically exact correlations and sampling in the two-dimensional Ising spin glass", en, Phys. Rev. E **87**, 043303 (2013).

- [193] A. Cavagna, T. S. Grigera, and P. Verrocchio, "Mosaic Multistate Scenario Versus One-State Description of Supercooled Liquids", Phys. Rev. Lett. **98** (2007).
- [194] G. Biroli, J.-P. Bouchaud, A. Cavagna, T. S. Grigera, and P. Verrocchio, "Thermodynamic signature of growing amorphous order in glass-forming liquids", Nat. Phys. **4**, 771–775 (2008).
- [195] G. M. Hocky, T. E. Markland, and D. R. Reichman, "Growing Point-to-Set Length Scale Correlates with Growing Relaxation Times in Model Supercooled Liquids", Phys. Rev. Lett. **108** (2012).
- [196] L. Berthier, P. Charbonneau, and S. Yaida, "Efficient measurement of point-to-set correlations and overlap fluctuations in glass-forming liquids", The Journal of Chemical Physics **144** (2016).
- [197] L. Berthier, P. Charbonneau, D. Coslovich, A. Ninarello, M. Ozawa, and S. Yaida, "Configurational entropy measurements in extremely supercooled liquids that break the glass ceiling", Proc. Natl. Acad. Sci. U.S.A. **114**, 11356–11361 (2017).
- [198] V. Lubchenko and P. G. Wolynes, "Theory of Structural Glasses and Supercooled Liquids", Annual Review of Physical Chemistry **58** (2007).
- [199] J.-P. Bouchaud, L. F. Cugliandolo, J. Kurchan, and M. Mézard, "Out of equilibrium dynamics in spin-glasses and other glassy systems", **12**, 161 (1998).
- [200] M. D. Ediger, "Spatially Heterogeneous Dynamics in Supercooled Liquids", Annual Review of Physical Chemistry **51**, 99–128 (2000).
- [201] S. C. Glotzer, "Spatially heterogeneous dynamics in liquids: insights from simulation", Journal of Non-Crystalline Solids **274**, 342–355 (2000).
- [202] L. Berthier, G. Biroli, J.-P. Bouchaud, L. Cipelletti, and W. van Saarloos, *Dynamical Heterogeneities in Glasses, Colloids, and Granular Media* (Oxford University Press, July 2011).
- [203] R. N. Bhatt and A. P. Young, "Numerical studies of Ising spin glasses in two, three, and four dimensions", Phys. Rev. B **37**, 5606–5614 (1988).
- [204] J.-S. Wang and R. H. Swendsen, "Low-temperature properties of the \pm J Ising spin glass in two dimensions", Phys. Rev. B **38**, 4840–4844 (1988).
- [205] K. Hukushima and K. Nemoto, "A Monte Carlo study on the spin dynamics of the 2D+or-J Ising spin glass model", J. Phys.: Condens. Matter **5**, 1389–1398 (1993).
- [206] C. K. Thomas, D. A. Huse, and A. A. Middleton, "Zero- and Low-Temperature Behavior of the Two-Dimensional ± J Ising Spin Glass", Phys. Rev. Lett. **107**, 047203 (2011).
- [207] F. Barahona, R. Maynard, R. Rammal, and J. P. Uhry, "Morphology of ground states of two-dimensional frustration model", J. Phys. A: Math. Gen. 15, 673 (1982).
- [208] F. Barahona, "On the computational complexity of Ising spin glass models", J. Phys. A: Math. Gen. **15**, 3241 (1982).
- [209] H. Takayama, K. Nemoto, and H. Matsukawa, "Monte Carlo study of the 2D \pm J Ising spin glass dynamical aspects of spin clusters", J. Magn. Magn. Mat. **31-34**, 1303–1304 (1983).
- [210] D. A. Huse, "Monte Carlo simulation study of domain growth in an Ising spin glass", Phys. Rev. B **43**, 8673–8675 (1991).
- [211] F. Romá, S. Bustingorry, and P. M. Gleiser, "Signature of the Ground-State Topology in the Low-Temperature Dynamics of Spin Glasses", Phys. Rev. Lett. **96**, 167205 (2006).

- [212] N. Xu, K.-H. Wu, S. J. Rubin, Y.-J. Kao, and A. W. Sandvik, "Dynamic scaling in the two-dimensional Ising spin glass with normal-distributed couplings", Phys. Rev. E **96**, 052102 (2017).
- [213] S. J. Rubin, N. Xu, and A. W. Sandvik, "Dual time scales in simulated annealing of a two-dimensional Ising spin glass", en, Phys. Rev. E **95**, 052133 (2017).
- [214] L. A. Fernández, E. Marinari, V. Martín-Mayor, G. Parisi, and J. J. Ruiz-Lorenzo, "Out-of-equilibrium 2D Ising spin glass: almost, but not quite, a free-field theory", en, J. Stat. Mech. **2018**, 103301 (2018).
- [215] H. Khoshbakht and M. Weigel, "Domain-wall excitations in the two-dimensional Ising spin glass", Phys. Rev. B **97**, 064410 (2018).
- [216] L. A. Fernández, E. Marinari, V. Martín-Mayor, I. Paga, and J. J. Ruiz-Lorenzo, "Dimensional crossover in the aging dynamics of spin glasses in a film geometry", en, Phys. Rev. B **100**, 184412 (2019).
- [217] L. A. Fernández, E. Marinari, V. Martín-Mayor, G. Parisi, and J. J. Ruiz-Lorenzo, "An experiment-oriented analysis of 2D spin-glass dynamics: a twelve time-decades scaling study", en, J. Phys. A: Math. Theor. **52**, 224002 (2019).
- [218] J. Kisker, L. Santen, M. Schreckenberg, and H. Rieger, "Off-equilibrium dynamics in finite-dimensional spin-glass models", Phys. Rev. B **53**, 6418–6428 (1996).
- [219] E. Vincent, V. Dupuis, M. Alba, J. Hammann, and J.-P. Bouchaud, "Aging phenomena in spin-glass and ferromagnetic phases: Domain growth and wall dynamics", Europhysics Letters (EPL) **50**, 674–680 (2000).
- [220] L. F. Cugliandolo, "Topics in coarsening phenomena", Physica A **389**, 4360–4373 (2010).
- [221] H. Rieger, "Monte Carlo Studies of Ising Spin Glasses and Random Field Systems", in *Annual reviews of computational physics*, Vol. 2, edited by D. Stauffer (World Scientific, 1995), p. 295.
- [222] M. Krasnytska, B. Berche, Y. Holovatch, and R. Kenna, "Ising model with variable spin/agent strengths", J. Phys.: Complexity 1, 035008 (2020).
- [223] M. Krasnytska, B. Berche, Y. Holovatch, and R. Kenna, "Generalized Ising Model on a Scale-Free Network: An Interplay of Power Laws", Entropy **23**, 1175 (2021).
- [224] M. Dudka, M. Krasnytska, J. J. Ruiz-Lorenzo, and Y. Holovatch, "Effective and asymptotic criticality of structurally disordered magnets", J. Magn. Magn. Mat. **575**, 170718 (2023).
- [225] M. Henkel and M. Pleimling, "Superuniversality in phase-ordering disordered ferromagnets", **78**, 224419 (2008).
- [226] A. Bray, "Theory of phase-ordering kinetics", Adv. Phys. 43, 357–459 (1994).
- [227] D. A. Huse and C. L. Henley, "Pinning and Roughening of Domain Walls in Ising Systems Due to Random Impurities", Phys. Rev. Lett. **54**, 2708–2711 (1985).
- [228] A. R. Rivera, G. A. Pérez Alcázar, and J. A. Plascak, "Diluted and random-bond Ising model: Application to the Fe-Mn-Al alloys", Phys. Rev. B **41**, 4774–4777 (1990).
- [229] A. Sicilia, J. J. Arenzon, A. J. Bray, and L. F. Cugliandolo, "Geometric properties of two-dimensional coarsening with weak disorder", EPL **82**, 10001 (2008).
- [230] J. L. Iguain, S. Bustingorry, A. B. Kolton, and L. F. Cugliandolo, "Growing correlations and aging of an elastic line in a random potential", Phys. Rev. B **80**, 094201 (2009).

- [231] F. Corberi, E. Lippiello, A. Mukherjee, S. Puri, and M. Zannetti, "Growth law and superuniversality in the coarsening of disordered ferromagnets", J. Stat. Mech. **2011**, P03016 (2011).
- [232] F. Corberi, L. F. Cugliandolo, F. Insalata, and M. Picco, "Coarsening and percolation in a disordered ferromagnet", Phys. Rev. E **95**, 022101 (2017).
- [233] R. Paul, S. Puri, and H. Rieger, "Domain growth in random magnets", EPL **68**, 881–887 (2004).
- [234] J. Charfreitag, M. Jünger, S. Mallach, and P. Mutzel, "McSparse: Exact Solutions of Sparse Maximum Cut and Sparse Unconstrained Binary Quadratic Optimization Problems", in 2022 proceedings of the symposium on algorithm engineering and experiments, edited by C. A. Phillips and B. Speckmann (2022), p. 54.
- [235] F. Romá, S. Risau-Gusman, A. J. Ramírez-Pastor, F. Nieto, and E. E. Vogel, "The ground state energy of the Edwards-Anderson spin glass model with a parallel tempering Monte Carlo algorithm", Physica A **388**, 2821 (2009).
- [236] W. Wang, J. Machta, and H. G. Katzgraber, "Comparing Monte Carlo methods for finding ground states of Ising spin glasses: Population annealing, simulated annealing, and parallel tempering", Phys. Rev. E **92**, 013303 (2015).
- [237] H. Khoshbakht and M. Weigel, "Domain-wall excitations in the two-dimensional Ising spin glass", Phys. Rev. B **97** (2018).
- [238] L. Münster, C. Norrenbrock, A. K. Hartmann, and A. P. Young, "Ordering behavior of the two-dimensional Ising spin glass with long-range correlated disorder", Phys. Rev. E **103** (2021).
- [239] M. Serbyn, Z. Papić, and D. A. Abanin, "Local Conservation Laws and the Structure of the Many-Body Localized States", Phys. Rev. Lett. **111** (2013).
- [240] D. A. Huse, R. Nandkishore, and V. Oganesyan, "Phenomenology of fully many-body-localized systems", Phys. Rev. B **90** (2014).
- [241] V. Ros, M. Müller, and A. Scardicchio, "Integrals of motion in the many-body localized phase", Nucl. Phys. B **891** (2015).
- [242] J. Z. Imbrie, V. Ros, and A. Scardicchio, "Local integrals of motion in many-body localized systems", Ann. Phys. **529** (2017).
- [243] W. De Roeck and F. Huveneers, "Stability and instability towards delocalization in many-body localization systems", en, Phys. Rev. B **95** (2017).
- [244] P. Sierant, M. Lewenstein, and J. Zakrzewski, "Polynomially Filtered Exact Diagonalization Approach to Many-Body Localization", en, Phys. Rev. Lett. **125** (2020).
- [245] P. Sierant, D. Delande, and J. Zakrzewski, "Thouless Time Analysis of Anderson and Many-Body Localization Transitions", en, Phys. Rev. Lett. **124** (2020).
- [246] D. J. Luitz, N. Laflorencie, and F. Alet, "Many-body localization edge in the random-field Heisenberg chain", en, Phys. Rev. B **91** (2015).
- [247] N. Macé, F. Alet, and N. Laflorencie, "Multifractal Scalings Across the Many-Body Localization Transition", en, Phys. Rev. Lett. **123** (2019).
- [248] D. J. Luitz, F. Huveneers, and W. De Roeck, "How a Small Quantum Bath Can Thermalize Long Localized Chains", en, Phys. Rev. Lett. **119** (2017).
- [249] T. Thiery, F. Huveneers, M. Müller, and W. De Roeck, "Many-Body Delocalization as a Quantum Avalanche", en, Phys. Rev. Lett. **121** (2018).

- [250] M. Goihl, J. Eisert, and C. Krumnow, "Exploration of the stability of many-body localized systems in the presence of a small bath", en, Phys. Rev. B **99** (2019).
- [251] P. J. D. Crowley and A. Chandran, "Avalanche induced coexisting localized and thermal regions in disordered chains", Phys. Rev. Res. **2** (2020).
- [252] J. Léonard, S. Kim, M. Rispoli, A. Lukin, R. Schittko, J. Kwan, E. Demler, D. Sels, and M. Greiner, "Probing the onset of quantum avalanches in a many-body localized system", en, Nat. Phys. **19** (2023).
- [253] J. C. Peacock and D. Sels, "Many-body delocalization from embedded thermal inclusion", en, Phys. Rev. B **108** (2023).
- [254] T. Szołdra, P. Sierant, M. Lewenstein, and J. Zakrzewski, "Catching thermal avalanches in the disordered XXZ model", en, Phys. Rev. B **109** (2024).
- [255] J. Šuntajs and L. Vidmar, "Ergodicity Breaking Transition in Zero Dimensions", Phys. Rev. Lett. **129** (2022).
- [256] J. Léonard, S. Kim, M. Rispoli, A. Lukin, R. Schittko, J. Kwan, E. Demler, D. Sels, and M. Greiner, "Probing the onset of quantum avalanches in a many-body localized system", Nat. Phys. **19**, 481–485 (2023).
- [257] A. Morningstar, L. Colmenarez, V. Khemani, D. J. Luitz, and D. A. Huse, "Avalanches and many-body resonances in many-body localized systems", Phys. Rev. B **105** (2022).
- [258] H. Ha, A. Morningstar, and D. A. Huse, "Many-Body Resonances in the Avalanche Instability of Many-Body Localization", Phys. Rev. Lett. **130** (2023).
- [259] K. Pawlik, P. Sierant, L. Vidmar, and J. Zakrzewski, "Many-body mobility edge in quantum sun models", Phys. Rev. B **109** (2024).
- [260] L. Colmenarez, D. J. Luitz, and W. De Roeck, "Ergodic inclusions in many-body localized systems", Phys. Rev. B **109** (2024).
- [261] K. Pawlik, N. Laflorencie, and J. Zakrzewski, *Unconventional Thermalization of a Localized Chain Interacting with an Ergodic Bath*, 2025.
- [262] F. Pietracaprina, N. Macé, D. J. Luitz, and F. Alet, "Shift-invert diagonalization of large many-body localizing spin chains", SciPost Phys. **5** (2018).
- [263] N. Laflorencie, G. Lemarié, and N. Macé, "Chain breaking and Kosterlitz-Thouless scaling at the many-body localization transition in the random-field Heisenberg spin chain", Phys. Rev. Res. **2** (2020).
- [264] J. Z. Imbrie, "Diagonalization and Many-Body Localization for a Disordered Quantum Spin Chain", en, Phys. Rev. Lett. **117** (2016).
- [265] B. Villalonga and B. K. Clark, 2020.
- [266] S. J. Garratt, S. Roy, and J. T. Chalker, "Local resonances and parametric level dynamics in the many-body localized phase", Phys. Rev. B **104** (2021).
- [267] P. Crowley and A. Chandran, "A constructive theory of the numerically accessible many-body localized to thermal crossover", SciPost Phys. **12** (2022).
- [268] S. Gopalakrishnan, M. Müller, V. Khemani, M. Knap, E. Demler, and D. A. Huse, "Low-frequency conductivity in many-body localized systems", Phys. Rev. B **92** (2015).
- [269] S. D. Geraedts, R. Nandkishore, and N. Regnault, "Many-body localization and thermalization: Insights from the entanglement spectrum", Phys. Rev. B **93** (2016).

- [270] V. Khemani, S. Lim, D. Sheng, and D. A. Huse, "Critical Properties of the Many-Body Localization Transition", Phys. Rev. X 7 (2017).
- [271] L. A. Colmenarez, P. A. McClarty, M. Haque, and D. J. Luitz, "Statistics of correlation functions in the random Heisenberg chain", SciPost Phys. 7 (2019).
- [272] B. Villalonga and B. K. Clark, 2020.
- [273] S. J. Garratt and S. Roy, "Resonant energy scales and local observables in the many-body localized phase", Phys. Rev. B **106** (2022).
- [274] J. A. Kjäll, "Many-body localization and level repulsion", Phys. Rev. B 97 (2018).
- [275] J. Z. Imbrie, "On Many-Body Localization for Quantum Spin Chains", J. Stat. Phys. **163**, 998–1048 (2016).
- [276] W. D. Roeck, L. Giacomin, F. Huveneers, and O. Prosniak, *Absence of Normal Heat Conduction in Strongly Disordered Interacting Quantum Chains*, 2025.
- [277] G. De Tomasi, I. M. Khaymovich, F. Pollmann, and S. Warzel, "Rare thermal bubbles at the many-body localization transition from the Fock space point of view", Phys. Rev. B **104** (2021).
- [278] D. M. Long, P. J. Crowley, V. Khemani, and A. Chandran, "Phenomenology of the Prethermal Many-Body Localized Regime", Phys. Rev. Lett. **131** (2023).
- [279] J. Colbois, F. Alet, and N. Laflorencie, "Interaction-Driven Instabilities in the Random-Field X X Z Chain", en, Phys. Rev. Lett. **133** (2024).
- [280] J. Colbois, F. Alet, and N. Laflorencie, "Statistics of systemwide correlations in the random-field XXZ chain: Importance of rare events in the many-body localized phase", en, Phys. Rev. B **110** (2024).
- [281] G. Biroli, A. K. Hartmann, and M. Tarzia, "Large-deviation analysis of rare resonances for the many-body localization transition", en, Phys. Rev. B **110** (2024).
- [282] A. Pal and D. A. Huse, "Many-body localization phase transition", en, Phys. Rev. B 82 (2010).
- [283] A. De Luca and A. Scardicchio, "Ergodicity breaking in a model showing many-body localization", EPL **101** (2013).
- [284] J. Gray, S. Bose, and A. Bayat, "Many-body localization transition: Schmidt gap, entanglement length, and scaling", Phys. Rev. B **97** (2018).
- [285] E. V. H. Doggen, F. Schindler, K. S. Tikhonov, A. D. Mirlin, T. Neupert, D. G. Polyakov, and I. V. Gornyi, "Many-body localization and delocalization in large quantum chains", en, Phys. Rev. B **98** (2018).
- [286] D. Abanin, J. Bardarson, G. De Tomasi, S. Gopalakrishnan, V. Khemani, S. Parameswaran, F. Pollmann, A. Potter, M. Serbyn, and R. Vasseur, "Distinguishing localization from chaos: Challenges in finite-size systems", en, Ann. Phys. (N. Y.) **427** (2021).
- [287] J.-L. Pichard, Random Transfer Matrix Theory and Conductance Fluctuations (Springer US, 1991).
- [288] P. Markoš, "Electron transport in strongly disordered structures", Phys. B: Condensed Matter **405** (2010).
- [289] G. Lemarié, "Glassy Properties of Anderson Localization: Pinning, Avalanches, and Chaos", en, Phys. Rev. Lett. **122** (2019).
- [290] G. Feher and E. A. Gere, "Electron Spin Resonance Experiments on Donors in Silicon. II. Electron Spin Relaxation Effects", Physical Review **114**, 1245–1256 (1959).

- [291] E. Abrahams, 50 Years of Anderson Localization (WORLD SCIENTIFIC, June 2010).
- [292] A. Scardicchio and T. Thiery, *Perturbation theory approaches to Anderson and Many-Body Localization: some lecture notes*, 2017.
- [293] P. Mörters, R. Moser, M. Penrose, H. Schwetlick, and J. Zimmer, eds., *Analysis and Stochastics of Growth Processes and Interface Models* (Oxford University Press, Oxford, UK, 2008).
- [294] P. W. Anderson, D. J. Thouless, E. Abrahams, and D. S. Fisher, "New method for a scaling theory of localization", Phys. Rev. B **22**, 3519–3526 (1980).
- [295] J. Šuntajs, T. Prosen, and L. Vidmar, "Localization challenges quantum chaos in the finite two-dimensional Anderson model", Phys. Rev. B **107** (2023).
- [296] E. Abrahams, P. W. Anderson, D. C. Licciardello, and T. V. Ramakrishnan, "Scaling Theory of Localization: Absence of Quantum Diffusion in Two Dimensions", Phys. Rev. Lett. **42**, 673–676 (1979).
- [297] B. I. Shklovskii, B. Shapiro, B. R. Sears, P. Lambrianides, and H. B. Shore, "Statistics of spectra of disordered systems near the metal-insulator transition", Phys. Rev. B 47, 11487–11490 (1993).
- [298] E. Hofstetter and M. Schreiber, "Statistical properties of the eigenvalue spectrum of the three-dimensional Anderson Hamiltonian", Phys. Rev. B **48**, 16979–16985 (1993).
- [299] I. K. Zharekeshev and B. Kramer, "Scaling of level statistics at the disorder-induced metal-insulator transition", Phys. Rev. B **51**, 17239–17242 (1995).
- [300] I. K. Zharekeshev and B. Kramer, "Asymptotics of Universal Probability of Neighboring Level Spacings at the Anderson Transition", Phys. Rev. Lett. **79**, 717–720 (1997).
- [301] I. Varga, E. Hofstetter, M. Schreiber, and J. Pipek, "Shape analysis of the level-spacing distribution around the metal-insulator transition in the three-dimensional Anderson model", Phys. Rev. B **52**, 7783–7786 (1995).
- [302] A. Rodriguez, L. J. Vasquez, K. Slevin, and R. A. Römer, "Critical Parameters from a Generalized Multifractal Analysis at the Anderson Transition", Phys. Rev. Lett. **105** (2010).
- [303] A. Rodriguez, L. J. Vasquez, K. Slevin, and R. A. Römer, "Multifractal finite-size scaling and universality at the Anderson transition", Phys. Rev. B **84** (2011).
- [304] K. Slevin and T. Ohtsuki, "Corrections to Scaling at the Anderson Transition", Phys. Rev. Lett. **82**, 382–385 (1999).
- [305] K. Slevin and T. Ohtsuki, "Critical exponent for the Anderson transition in the three-dimensional orthogonal universality class", New Journal of Physics **16**, 015012 (2014).
- [306] Y. Ueoka and K. Slevin, "Dimensional Dependence of Critical Exponent of the Anderson Transition in the Orthogonal Universality Class", Journal of the Physical Society of Japan 83, 084711 (2014).
- [307] E. Tarquini, G. Biroli, and M. Tarzia, "Critical properties of the Anderson localization transition and the high-dimensional limit", Phys. Rev. B **95** (2017).
- [308] G. Biroli, G. Semerjian, and M. Tarzia, "Anderson Model on Bethe Lattices: Density of States, Localization Properties and Isolated Eigenvalue", Progress of Theoretical Physics Supplement **184**, 187–199 (2010).

- [309] F. Evers and A. D. Mirlin, "Anderson transitions", Rev. Mod. Phys. **80**, 1355–1417 (2008).
- [310] A. D. Mirlin and Y. V. Fyodorov, "Distribution of local densities of states, order parameter function, and critical behavior near the Anderson transition", Phys. Rev. Lett. **72**, 526–529 (1994).
- [311] A. D. Mirlin and Y. V. Fyodorov, "Statistical properties of one-point Green functions in disordered systems and critical behavior near the Anderson transition", Journal de Physique I **4**, 655–673 (1994).
- [312] A. M. García-García and E. Cuevas, "Dimensional dependence of the metal-insulator transition", Phys. Rev. B **75** (2007).
- [313] C. Castellani, C. D. Castro, and L. Peliti, "On the upper critical dimension in Anderson localisation", Journal of Physics A: Mathematical and General **19**, L1099–L1103 (1986).
- [314] A. D. Mirlin and F. Evers, "Multifractality and critical fluctuations at the Anderson transition", Phys. Rev. B **62**, 7920–7933 (2000).
- [315] F. Evers and A. D. Mirlin, "Fluctuations of the Inverse Participation Ratio at the Anderson Transition", Phys. Rev. Lett. **84**, 3690–3693 (2000).
- [316] F. Wegner, "Inverse participation ratio in $2+\varepsilon$ dimensions", Zeitschrift für Physik B Condensed Matter and Quanta **36**, 209–214 (1980).
- [317] K. Ishii, "Localization of Eigenstates and Transport Phenomena in the One-Dimensional Disordered System", Progress of Theoretical Physics Supplement **53**, 77–138 (1973).
- [318] J. C. Kimball, "Localisation and spectra in solid state systems", Journal of Physics C: Solid State Physics 11, 4347–4354 (1978).
- [319] E. N. Economou, *Green's Functions in Quantum Physics* (Springer Berlin Heidelberg, 2006).
- [320] N. Mott, "Conduction in glasses containing transition metal ions", Journal of Non-Crystalline Solids 1, 1–17 (1968).
- [321] L. Fleishman and P. W. Anderson, "Interactions and the Anderson transition", Phys. Rev. B **21**, 2366–2377 (1980).
- [322] F. Pietracaprina, V. Ros, and A. Scardicchio, "Forward approximation as a mean-field approximation for the Anderson and many-body localization transitions", en, Phys. Rev. B **93** (2016).
- [323] G. Parisi, S. Pascazio, F. Pietracaprina, V. Ros, and A. Scardicchio, "Anderson transition on the Bethe lattice: an approach with real energies", J. Phys. A: Math. Theor. 53, 014003 (2020).
- [324] K. Agarwal, S. Gopalakrishnan, M. Knap, M. Müller, and E. Demler, "Anomalous Diffusion and Griffiths Effects Near the Many-Body Localization Transition", Phys. Rev. Lett. **114** (2015).
- [325] S. Gopalakrishnan and S. Parameswaran, "Dynamics and transport at the threshold of many-body localization", Physics Reports **862**, 1–62 (2020).
- [326] P. Bordia, H. Lüschen, U. Schneider, M. Knap, and I. Bloch, "Periodically driving a many-body localized quantum system", Nat. Phys. **13**, 460–464 (2017).
- [327] J.-y. Choi, S. Hild, J. Zeiher, P. Schauß, A. Rubio-Abadal, T. Yefsah, V. Khemani, D. A. Huse, I. Bloch, and C. Gross, "Exploring the many-body localization transition in two dimensions", Science **352**, 1547–1552 (2016).

- [328] L. F. Cugliandolo and M. Mueller, "Quantum Glasses a Review", arXiv preprint arXiv:2208.05417 (2022).
- [329] K. Binder and A. P. Young, "Spin glasses: Experimental facts, theoretical concepts, and open questions", Rev. Mod. Phys. **58**, 801 (1986).
- [330] G. Biroli, D. Facoetti, M. Schiró, M. Tarzia, and P. Vivo, "Out-of-equilibrium phase diagram of the quantum random energy model", Phys. Rev. B **103** (2021).
- [331] 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** (2016).
- [332] Y. Y. Goldschmidt, "Solvable model of the quantum spin glass in a transverse field", Phys. Rev. B **41**, 4858–4861 (1990).
- [333] T. Jörg, F. Krzakala, J. Kurchan, and A. C. Maggs, "Simple Glass Models and Their Quantum Annealing", Phys. Rev. Lett. **101** (2008).
- [334] M. Rigol, V. Dunjko, and M. Olshanii, "Thermalization and its mechanism for generic isolated quantum systems", Nature **452**, 854–858 (2008).
- [335] M. Rigol, V. Dunjko, V. Yurovsky, and M. Olshanii, "Relaxation in a Completely Integrable Many-Body Quantum System: AnAb InitioStudy of the Dynamics of the Highly Excited States of 1D Lattice Hard-Core Bosons", Phys. Rev. Lett. **98** (2007).
- [336] J. Eisert, M. Friesdorf, and C. Gogolin, "Quantum many-body systems out of equilibrium", Nat. Phys. 11, 124–130 (2015).
- [337] C. J. Turner, A. A. Michailidis, D. A. Abanin, M. Serbyn, and Z. Papić, "Weak ergodicity breaking from quantum many-body scars", Nat. Phys. **14**, 745–749 (2018).
- [338] M. Serbyn, D. A. Abanin, and Z. Papić, "Quantum many-body scars and weak breaking of ergodicity", Nat. Phys. **17**, 675–685 (2021).
- [339] P. Sala, T. Rakovszky, R. Verresen, M. Knap, and F. Pollmann, "Ergodicity Breaking Arising from Hilbert Space Fragmentation in Dipole-Conserving Hamiltonians", Physical Review X **10** (2020).
- [340] S. Moudgalya, B. A. Bernevig, and N. Regnault, "Quantum many-body scars and Hilbert space fragmentation: a review of exact results", Reports on Progress in Physics 85, 086501 (2022).
- [341] S. Weigert, "The problem of quantum integrability", Physica D: Nonlinear Phenomena **56**, 107–119 (1992).
- [342] J.-S. Caux and J. Mossel, "Remarks on the notion of quantum integrability", Journal of Statistical Mechanics: Theory and Experiment **2011**, P02023 (2011).
- [343] V. E. Korepin, N. M. Bogoliubov, and A. G. Izergin, *Quantum Inverse Scattering Method and Correlation Functions* (Cambridge University Press, 1993).
- [344] E. Ilievski, M. Medenjak, T. Prosen, and L. Zadnik, "Quasilocal charges in integrable lattice systems", Journal of Statistical Mechanics: Theory and Experiment **2016**, 064008 (2016).
- [345] M. Jimbo and T. Miwa, "Algebraic analysis of solvable lattice models", American Mathematical Society Translations Series 2 **85** (1995).
- [346] F. H. L. Essler and M. Fagotti, "Quench dynamics and relaxation in isolated integrable quantum spin chains", Journal of Statistical Mechanics: Theory and Experiment **2016**, 064002 (2016).

- [347] J. Durnin, M. J. Bhaseen, and B. Doyon, "Nonequilibrium Dynamics and Weakly Broken Integrability", Phys. Rev. Lett. **127**, 130601 (2021).
- [348] S.-H. Lin, B. Sbierski, F. Dorfner, C. Karrasch, and F. Heidrich-Meisner, "Many-body localization of spinless fermions with attractive interactions in one dimension", SciPost Phys. 4, 002 (2018).
- [349] B. L. Altshuler, Y. Gefen, A. Kamenev, and L. S. Levitov, "Quasiparticle Lifetime in a Finite System: A Nonperturbative Approach", Phys. Rev. Lett. **78**, 2803–2806 (1997).
- [350] G. Biroli and M. Tarzia, "Delocalized glassy dynamics and many-body localization", Phys. Rev. B **96** (2017).
- [351] D. E. Logan and S. Welsh, "Many-body localization in Fock space: A local perspective", Phys. Rev. B **99** (2019).
- [352] K. Tikhonov and A. Mirlin, "From Anderson localization on random regular graphs to many-body localization", Ann. Phys. (N. Y.) **435**, 168525 (2021).
- [353] S. Roy and D. E. Logan, "The Fock-space landscape of many-body localisation", J. Phys. Condens. Matter **37** (2024).
- [354] K. S. Tikhonov, A. D. Mirlin, and M. A. Skvortsov, "Anderson localization and ergodicity on random regular graphs", Phys. Rev. B **94** (2016).
- [355] K. S. Tikhonov and A. D. Mirlin, "Critical behavior at the localization transition on random regular graphs", Phys. Rev. B **99** (2019).
- [356] K. S. Tikhonov and A. D. Mirlin, "Statistics of eigenstates near the localization transition on random regular graphs", Phys. Rev. B **99** (2019).
- [357] P. Prelov šek, O. S. Bari ši ć, and M. Mierzejewski, "Reduced-basis approach to many-body localization", Phys. Rev. B **97**, 035104 (2018).
- [358] S. J. Thomson and M. Schiró, "Time evolution of many-body localized systems with the flow equation approach", Phys. Rev. B **97**, 060201 (2018).
- [359] G. De Tomasi, F. Pollmann, and M. Heyl, "Efficiently solving the dynamics of many-body localized systems at strong disorder", Phys. Rev. B **99**, 241114 (2019).
- [360] J. Colbois and N. Laflorencie, "Breaking the chains: Extreme value statistics and localization in random spin chains", Phys. Rev. B **108**, 144206 (2023).
- [361] 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** (2015).
- [362] V. Oganesyan and D. A. Huse, "Localization of interacting fermions at high temperature", Phys. Rev. B **75** (2007).
- [363] M. V. Berry and M. Tabor, "Level clustering in the regular spectrum", Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences **356**, 375–394 (1977).
- [364] O. Bohigas, M. J. Giannoni, and C. Schmit, "Characterization of chaotic quantum spectra and universality of level fluctuation laws", Phys. Rev. Lett. **52**, 1–4 (1984).
- [365] D. Poilblanc, T. Ziman, J. Bellissard, F. Mila, and G. Montambaux, "Poisson vs GOE statistics in integrable and non-integrable quantum Hamiltonians", Europhysics Letters (EPL) **22**, 537 (1993).

- [366] R. Berkovits and B. I. Shklovskii, "Statistics of energy spectra of strongly disordered interacting electrons", Journal of Physics: Condensed Matter 11, 779–? (1999).
- [367] Y. Avishai, J. Richert, and R. Berkovits, "Level statistics in a Heisenberg chain with random magnetic field", Phys. Rev. B **66**, 052416 (2002).
- [368] E. S. Sørensen, M.-S. Chang, N. Laflorencie, and I. Affleck, "Quantum impurity entanglement", Journal of Statistical Mechanics: Theory and Experiment **2007**, P08003–P08003 (2007).
- [369] V. Alba, M. Fagotti, and P. Calabrese, "Entanglement entropy of excited states", Journal of Statistical Mechanics: Theory and Experiment **2009**, P10020 (2009).
- [370] J. Sato, B. Aufgebauer, H. Boos, F. Göhmann, A. Klümper, M. Takahashi, and C. Trippe, "Computation of Static Heisenberg-Chain Correlators: Control over Length and Temperature Dependence", Phys. Rev. Lett. **106** (2011).
- [371] J. P. Keating, N. Linden, and H. J. Wells, "Spectra and Eigenstates of Spin Chain Hamiltonians", Communications in Mathematical Physics **338**, 81–102 (2015).
- [372] V. Alba, "Eigenstate thermalization hypothesis and integrability in quantum spin chains", Phys. Rev. B **91** (2015).
- [373] 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**, P09005 (2013).
- [374] J. A. Kjäll, J. H. Bardarson, and F. Pollmann, "Many-Body Localization in a Disordered Quantum Ising Chain", en, Phys. Rev. Lett. **113**, 107204 (2014).
- [375] R. Vosk, D. A. Huse, and E. Altman, "Theory of the Many-Body Localization Transition in One-Dimensional Systems", Physical Review X **5** (2015).
- [376] A. C. Potter, R. Vasseur, and S. Parameswaran, "Universal Properties of Many-Body Delocalization Transitions", Physical Review X 5 (2015).
- [377] X. Yu, D. J. Luitz, and B. K. Clark, "Bimodal entanglement entropy distribution in the many-body localization transition", Phys. Rev. B **94** (2016).
- [378] D. A. Huse, R. Nandkishore, V. Oganesyan, A. Pal, and S. L. Sondhi, "Localization-protected quantum order", Phys. Rev. B **88**, 014206 (2013).
- [379] D. Pekker, G. Refael, E. Altman, E. A. Demler, and V. Oganesyan, "Hilbert-glass transition: New universality of temperature-tuned many-body dynamical quantum criticality", Physical Review X 4, 011052 (2014).
- [380] R. Vasseur, A. J. Friedman, S. A. Parameswaran, and A. C. Potter, "Quantum criticality of hot random spin chains", Phys. Rev. B **93**, 134207 (2016).
- [381] C. Monthus, "Many-body localization: strong disorder renormalization group for the dynamics of Floquet systems", Journal of Statistical Mechanics: Theory and Experiment **2016**, 033113 (2016).
- [382] Y. Y. Atas and E. Bogomolny, "Multifractality of eigenfunctions in spin chains", Phys. Rev. E **86** (2012).
- [383] D. J. Luitz, F. Alet, and N. Laflorencie, "Universal Behavior beyond Multifractality in Quantum Many-Body Systems", Phys. Rev. Lett. **112** (2014).
- [384] M. Serbyn and J. E. Moore, "Spectral statistics across the many-body localization transition", Phys. Rev. B **93** (2016).

- [385] D. J. Luitz, N. Laflorencie, and F. Alet, "Extended slow dynamical regime close to the many-body localization transition", Phys. Rev. B **93** (2016).
- [386] P. Calabrese and J. Cardy, "Evolution of entanglement entropy in one-dimensional systems", J. Stat. Mech. **0504**, P04010 (2005).
- [387] P. Calabrese and J. Cardy, "Entanglement entropy and conformal field theory", Journal of Physics A: Mathematical and Theoretical **42**, 504005 (2009).
- [388] C. T. Asplund, A. Bernamonti, F. Galli, and T. Hartman, "Entanglement scrambling in 2d conformal field theory", Journal of High Energy Physics 9, 110 (2015).
- [389] V. Hubeny, M. Rangamani, and T. Takayanagi, "A covariant holographic entanglement entropy proposal", Journal of High Energy Physics **2007**, 062 (2007).
- [390] J. Aparicio, J. Abajo-Arrastia, and E. López, "Holographic evolution of entanglement entropy", Journal of High Energy Physics, 1–27 (2010).
- [391] T. Hartman and J. Maldacena, "Time evolution of entanglement entropy from black hole interiors", Journal of High Energy Physics, 1–28 (2013).
- [392] H. Liu and S. J. Suh, "Entanglement tsunami: Universal scaling in holographic thermalization", Phys. Rev. Lett. **112**, 011601 (2014).
- [393] H. Liu and S. J. Suh, "Entanglement growth during thermalization in holographic systems", Phys. Rev. D **89**, 066012 (2014).
- [394] H. Casini, H. Liu, and M. Mezei, Spread of entanglement and causality.
- [395] M. Fagotti and P. Calabrese, "Evolution of entanglement entropy following a quantum quench: Analytic results for the XY chain in a transverse magnetic field", Phys. Rev. A **78**, 010306 (2008).
- [396] I. Peschel and V. Eisler, "Reduced density matrices and entanglement entropy in free lattice models", Journal of Physics A: Mathematical and Theoretical **42**, 504003 (2009).
- [397] A. S. Buyskikh, M. Fagotti, J. Schachenmayer, F. Essler, and A. J. Daley, "Entanglement growth and correlation spreading with variable-range interactions in spin and fermionic tunneling models", Phys. Rev. A 93, 053620 (2016).
- [398] A. M. Lauchli and C. Kollath, "Spreading of correlations and entanglement after a quench in the one-dimensional Bose Hubbard model", Journal of Statistical Mechanics: Theory and Experiment, P05018 (2008).
- [399] H. Kim and D. A. Huse, "Ballistic spreading of entanglement in a diffusive nonintegrable system", Phys. Rev. Lett. **111**, 127205 (2013).
- [400] J. H. Bardarson, F. Pollmann, and J. E. Moore, "Unbounded Growth of Entanglement in Models of Many-Body Localization", Phys. Rev. Lett. **109** (2012).
- [401] M. Serbyn, Z. Papić, and D. A. Abanin, "Universal slow growth of entanglement in interacting strongly disordered systems", Phys. Rev. Lett. **110**, 260601 (2013).
- [402] W. W. Ho and D. A. Abanin, "Entanglement dynamics in quantum many-body systems", Phys. Rev. B **95**, 094302 (2017).
- [403] G. D. Chiara, S. Montangero, P. Calabrese, and R. Fazio, "Entanglement entropy dynamics of Heisenberg chains", Journal of Statistical Mechanics: Theory and Experiment **2006**, P03001–P03001 (2006).
- [404] T. Chanda, P. Sierant, and J. Zakrzewski, "Time dynamics with matrix product states: Many-body localization transition of large systems revisited", Phys. Rev. B **101** (2020).

- [405] F. Evers, I. Modak, and S. Bera, "Internal clock of many-body delocalization", Phys. Rev. B **108** (2023).
- [406] J. Richter and A. Pal, "Many-body localization and delocalization dynamics in the thermodynamic limit", Phys. Rev. B **105** (2022).
- [407] B. Derrida, "Random-energy model: An exactly solvable model of disordered systems", en, Phys. Rev. B **24** (1981).
- [408] B. Derrida, "A generalization of the Random Energy Model which includes correlations between energies", J. Phys. (Paris), Lett. **46** (1985).
- [409] M. Mezard and G. P. Virasoro; *Spin Glass Theory and Beyond : An Introduction to the Replica Method and Its Applications* (World Scientific Publishing Company, 1987).
- [410] B Derrida and E Gardner, "Solution of the generalised random energy model", J. Phys. C: Solid State Phys. **19** (1986).
- [411] D. Carpentier and P. Le Doussal, "Glass transition of a particle in a random potential, front selection in nonlinear renormalization group, and entropic phenomena in Liouville and sinh-Gordon models", en, Phys. Rev. E **63** (2001).
- [412] P. Charbonneau, E. Marinari, M. Mézard, G. Parisi, F. Ricci-Tersenghi, G. Sicuro, and F. Zamponi, *Spin Glass Theory and Far Beyond: Replica Symmetry Breaking After 40 Years* (World Scientific, 2023).
- [413] S. Datta, Cambridge studies in semiconductor physics and microelectronic engineering: Electronic transport in mesoscopic systems series number 3 (Cambridge University Press, Cambridge, England).
- [414] D. S. Fisher and P. A. Lee, "Relation between conductivity and transmission matrix", Phys. Rev. B **23** (1981).
- [415] A. De Luca, B. Altshuler, V. Kravtsov, and A. Scardicchio, "Anderson Localization on the Bethe Lattice: Nonergodicity of Extended States", en, Phys. Rev. Lett. 113 (2014).
- [416] J.-N. Herre, J. F. Karcher, K. S. Tikhonov, and A. D. Mirlin, "Ergodicity-to-localization transition on random regular graphs with large connectivity and in many-body quantum dots", Phys. Rev. B **108** (2023).
- [417] T. Scoquart, I. V. Gornyi, and A. D. Mirlin, "Role of Fock-space correlations in many-body localization", Phys. Rev. B **109** (2024).
- [418] R. Abou-Chacra, D. J. Thouless, and P. W. Anderson, "A selfconsistent theory of localization", Journal of Physics C: Solid State Physics **6**, 1734–1752 (1973).
- [419] G. Biroli and M. Tarzia, "Anomalous dynamics on the ergodic side of the many-body localization transition and the glassy phase of directed polymers in random media", Phys. Rev. B **102** (2020).
- [420] M. Tarzia, in preparation.
- [421] B. Derrida and H. Spohn, "Polymers on disordered trees, spin glasses, and traveling waves", J. Stat. Phys. **51** (1988).
- [422] B. Derrida and R. B. Griffiths, "Directed Polymers on Disordered Hierarchical Lattices", EPL 8 (1989).
- [423] B. Derrida, "Mean Field Theory of Directed Polymers in a Random Medium and Beyond", Physica Scripta **T38** (1991).

- [424] M. Mezard, G. Parisi, and M. Virasoro, *Spin Glass Theory and Beyond: An Introduction to the Replica Method and Its Applications* (World Scientific, 1986).
- [425] F. Comets, Directed Polymers in Random Environments (Springer, 2017).
- [426] Y. Shiferaw and Y. Y. Goldschmidt, "Localization of a polymer in random media: Relation to the localization of a quantum particle", Phys. Rev. E **63** (2001).
- [427] C. Monthus and T. Garel, "Anderson localization on the Cayley tree: multifractal statistics of the transmission at criticality and off criticality", J. Phys. A: Math. Theor. 44 (2011).
- [428] S. Mu, J. Gong, and G. Lemarié, "Kardar-Parisi-Zhang Physics in the Density Fluctuations of Localized Two-Dimensional Wave Packets", Phys. Rev. Lett. **132** (2024).
- [429] O. Dimitrova and M. Mézard, "The cavity method for quantum disordered systems: from transverse random field ferromagnets to directed polymers in random media", J. Stat. Mech. **2011** (2011).
- [430] C. Monthus and T. Garel, "Random transverse field Ising model on the Cayley tree: analysis via boundary strong disorder renormalization", J. Stat. Mech. **2012** (2012).
- [431] A. M. Somoza, P. Le Doussal, and M. Ortuño, "Unbinding transition in semi-infinite two-dimensional localized systems", en, Phys. Rev. B **91** (2015).
- [432] E. Tarquini, G. Biroli, and M. Tarzia, "Level Statistics and Localization Transitions of Lévy Matrices", en, Phys. Rev. Lett. **116** (2016).
- [433] B. Derrida, "Directed polymers in a random medium", Physica A: Stat. Mech. Appl. **163** (1990).
- [434] M. R. Evans and B. Derrida, "Improved bounds for the transition temperature of directed polymers in a finite-dimensional random medium", J. Stat. Phys. **69** (1992).
- [435] M. Aizenman and S. Molchanov, "Localization at large disorder and at extreme energies: An elementary derivations", Comm. Math. Phys. **157** (1993).
- [436] M. Aizenman and S. Warzel, "Resonant delocalization for random Schrödinger operators on tree graphs", J. Eur. Math. Soc. **15**, 1167–1222 (2013).
- [437] Y. V. Fyodorov and J.-P. Bouchaud, "Freezing and extreme-value statistics in a random energy model with logarithmically correlated potential", J. Phys. A: Math. Theor. **41** (2008).
- [438] N. Laflorencie, J. Colbois, and F. Alet, 2025.
- [439] E. Gardner and B. Derrida, "The probability distribution of the partition function of the random energy model", J. Phys. A: Math. Gen. **22** (1989).
- [440] A. Pieper, M. Kreutzer, A. Alvermann, M. Galgon, H. Fehske, G. Hager, B. Lang, and G. Wellein, "High-performance implementation of Chebyshev filter diagonalization for interior eigenvalue computations", en, J. Comput. Phys. **325** (2016).
- [441] A. Andreanov, M. Carrega, J. Murugan, J. Olle, D. Rosa, and R. Shir, "From Dyson models to many-body quantum chaos", Phys. Rev. B **111** (2025).
- [442] A. Abbout, G. Lemarié, and J.-L. Pichard, "Thermal Enhancement of Interference Effects in Quantum Point Contacts", Phys. Rev. Lett. **106** (2011).
- [443] C. Gorini, R. A. Jalabert, W. Szewc, S. Tomsovic, and D. Weinmann, "Theory of scanning gate microscopy", Phys. Rev. B **88** (2013).

- [444] C. Vanoni, B. L. Altshuler, V. E. Kravtsov, and A. Scardicchio, "Renormalization group analysis of the anderson model on random regular graphs", Proceedings of the National Academy of Sciences **121**, e2401955121 (2024).
- [445] J. Niedda, G. B. Testasecca, G. Magnifico, F. Balducci, C. Vanoni, and A. Scardicchio, "Renormalization group analysis of the many-body localization transition in the random-field xxz chain", Phys. Rev. B 112, 144201 (2025).
- [446] M. Körner, H. G. Katzgraber, and A. K. Hartmann, "Probing tails of energy distributions using importance-sampling in the disorder with a guiding function", J. Stat. Mech. **2006** (2006).
- [447] P. Sierant, M. Lewenstein, A. Scardicchio, and J. Zakrzewski, "Stability of many-body localization in Floquet systems", Phys. Rev. B **107** (2023).
- [448] T. Szołdra, P. Sierant, M. Lewenstein, and J. Zakrzewski, "Catching thermal avalanches in the disordered XXZ model", Phys. Rev. B **109** (2024).
- [449] P. Bordia, H. Lüschen, S. Scherg, S. Gopalakrishnan, M. Knap, U. Schneider, and I. Bloch, "Probing Slow Relaxation and Many-Body Localization in Two-Dimensional Quasiperiodic Systems", Phys. Rev. X 7 (2017).
- [450] H. P. Lüschen, P. Bordia, S. Scherg, F. Alet, E. Altman, U. Schneider, and I. Bloch, "Observation of Slow Dynamics near the Many-Body Localization Transition in One-Dimensional Quasiperiodic Systems", Phys. Rev. Lett. **119** (2017).
- [451] P. R. Amestoy, I. S. Duff, J.-Y. L'Excellent, and J. Koster, "A Fully Asynchronous Multifrontal Solver Using Distributed Dynamic Scheduling", SIAM J. Matrix Anal. Appl. **23** (2001).
- [452] P. R. Amestoy, A. Guermouche, J.-Y. L'Excellent, and S. Pralet, "Hybrid scheduling for the parallel solution of linear systems", Parallel Comput. **32** (2006).
- [453] O. Schenk and K. Gärtner, "Solving unsymmetric sparse systems of linear equations with PARDISO", Future Gener. Comput. Syst. **20** (2004).
- [454] M. Tarzia and M. A. Moore, "Glass phenomenology from the connection to spin glasses", Phys. Rev. E **75** (2007).

Sujet : Rupture d'ergodicité dans la matière désordonnée : Nouvelles méthodes à travers des systèmes physiques

Résumé: L'hypothèse ergodique se situe au cœur de la mécanique statistique, mais de nombreux systèmes physiques s'en écartent lorsqu'ils sont hors d'équilibre. Les mécanismes à l'origine de cette brisure d'ergodicité donnent naissance à des phénomènes physiques riches, mais posent également des défis considérables, en particulier pour les simulations numériques. Au fil des années, diverses méthodes ont été développées pour relever ces défis dans des contextes spécifiques. Cette thèse s'appuie sur ces approches établies et les adapte afin d'étudier des problèmes dans des systèmes physiques où de telles techniques ne sont pas encore courantes. Ce faisant, nous cherchons à mettre en évidence de nouvelles perspectives, tant sur les méthodes elles-mêmes que sur les systèmes désordonnés que nous examinons. Notre travail se concentre sur trois problèmes principaux dans l'étude plus large de la rupture d'ergodicité dans les systèmes désordonnés: Localisation à plusieurs corps, rupture spontanée de symétrie dans les systèmes de spins frustrés et les verres de spin.

Mots clés : rupture d'ergodicité ; transitions de phase ; verres de spin ; algorithme swap ; localisation à plusieurs corps ; grandes déviations, simulations Monte Carlo.

Subject: Ergodicity breaking in disordered matter: Novel methods across physical systems

Abstract: The ergodic hypothesis lies at the heart of statistical mechanics, yet many physical systems deviate from it under out-of-equilibrium conditions. The mechanisms behind this ergodicity breaking give rise to rich physical phenomena but also pose significant challenges, particularly for numerical simulations. Over the years, a variety of methods have been developed to address these challenges in specific contexts. This thesis builds on these established approaches, adapting them to investigate problems in physical systems where such techniques are not yet standard. By doing so, we aim to uncover new insights into both the methods themselves and the disordered systems we are studying. Our work focuses on three main problems within the broader study of ergodicity breaking in disordered systems: Many-body localization, spontaneous symmetry breaking in frustrated spin systems and spin glasses.

Keywords: ergodicity breaking; phase transitions; spin glasses; swap algorithm; manybody localization; large deviations, Monte Carlo simulations.