

Many-Body Localization and Thermalization in Quantum Statistical Mechanics

Rahul Nandkishore¹ and David A. Huse^{1,2}

¹Princeton Center for Theoretical Science, Princeton University, Princeton, New Jersey 08544; email: rahuln@princeton.edu, huse@princeton.edu

²Department of Physics, Princeton University, Princeton, New Jersey 08544

Annu. Rev. Condens. Matter Phys. 2015. 6:15–38

The *Annual Review of Condensed Matter Physics* is online at conmatphys.annualreviews.org

This article's doi:
10.1146/annurev-conmatphys-031214-014726

Copyright © 2015 by Annual Reviews.
All rights reserved

Keywords

closed systems, entanglement, eigenstate, nonequilibrium, glass

Abstract

We review some recent developments in the statistical mechanics of isolated quantum systems. We provide a brief introduction to quantum thermalization, paying particular attention to the eigenstate thermalization hypothesis (ETH) and the resulting single-eigenstate statistical mechanics. We then focus on a class of systems that fail to quantum thermalize and whose eigenstates violate the ETH: These are the many-body Anderson-localized systems; their long-time properties are not captured by the conventional ensembles of quantum statistical mechanics. These systems can forever locally remember information about their local initial conditions and are thus of interest for possibilities of storing quantum information. We discuss key features of many-body localization (MBL) and review a phenomenology of the MBL phase. Single-eigenstate statistical mechanics within the MBL phase reveal dynamically stable ordered phases, and phase transitions among them, that are invisible to equilibrium statistical mechanics and can occur at high energy and low spatial dimensionality, where equilibrium ordering is forbidden.

1. INTRODUCTION

Although many of the fundamentals of quantum statistical mechanics were formulated together with those of quantum mechanics itself almost a century ago (1), the subject has had a recent rebirth. This has occurred because of the development of methods and tools in atomic, molecular, optical, and condensed matter physics that allow one to build, control, and study in the laboratory many new sorts of assemblies of strongly interacting quantum degrees of freedom (2–7). Such systems extend our abilities to explore and understand many-body quantum mechanics and are also of interest because of the possibilities of using them as components of new quantum technologies.

Many formulations of quantum statistical mechanics postulate that the system of interest is in contact with an external reservoir (8), with certain properties of the reservoir often taken for granted. However, recent experimental progress in well-approximating isolated many-body quantum systems [such as cold gases of neutral atoms] motivates a fresh consideration of the statistical mechanics of closed quantum systems, i.e., isolated quantum systems not coupled to any external reservoirs (see, e.g., Reference 9 for a recent review with some overlap with the present review). The statistical mechanical description of closed quantum systems is also important as a point of principle because if we assume that any external reservoir or measuring apparatus is itself a quantum system, it can then be included as part of the closed many-body quantum system of interest.

One fundamental question about such a closed quantum many-body system is: What states does its unitary time evolution bring it to after an arbitrarily long time? To make the possible answers to this question sharp, one needs to consider the thermodynamic limit of a large system, as we discuss below. There appear to be two answers to this question that are robust under small but arbitrary local perturbations to the system's Hamiltonian, namely thermalization and localization. The answer can depend on the nature of the system and on the initial state being considered, and the system can show quantum phase transitions between these two possibilities, as the system or the initial state is varied. The main goal of this review is to define, discuss, and elaborate these two possibilities, reviewing some of what is either known or (mostly) conjectured about them.

In order for a closed quantum system to thermally equilibrate under its own dynamics, the system must be able to act as its own reservoir, i.e., the dynamics must be such that for a subsystem that contains only a small fraction of the degrees of freedom of the entire system, the coupling to the rest of the system mimics a coupling to a reservoir. If the dynamics satisfy this property, the microcanonical, canonical, and grand canonical ensembles for the full system all give the correct long-time equilibrium properties of subsystems; the conventional theory of quantum statistical mechanics thus applies to the long-time steady states of subsystems, and we say the system thermalizes. However, not all closed quantum systems act as reservoirs that thermalize their subsystems.

Localized systems, first identified by Anderson (10), do not act as reservoirs for themselves, and thus do not thermalize. Instead, the long-time states of subsystems are determined by (and thus can remember) some local details of the system's initial state. This is why localized systems are of interest as possible quantum memories. The distinction between localization and thermalization is only dynamical. It is invisible if one examines only thermodynamic quantities, which are determined by averaging over an equilibrium ensemble of states, and essentially assumes thermalization. However, the distinction is quite apparent if one looks at the properties of individual exact many-body eigenstates of the system's Hamiltonian. Thus, we examine the single-eigenstate limit of the microcanonical ensemble, which is able to detect the quantum phase transition (or transitions) between the localized and thermalizing phases. This suggests a new eigenstate statistical mechanics, which is very useful in investigating localization, and also reveals a whole new world of localized phases, and quantum phase transitions between them, within the localized regime (11).

The bulk of the existing literature about localization focuses on noninteracting systems or on the low temperature limit. However, the perturbative arguments found in References 12–14 and particularly in Reference 15, numerical exact diagonalization studies, such as those in References 16 and 17, and even a recent mathematical proof (18) have provided strong evidence that localization can occur in highly excited states of strongly interacting many-body quantum systems—a phenomenon that has now come to be known as many-body localization (MBL) and has been the subject of considerable recent work (19–45).

MBL represents a new frontier of quantum statistical mechanics. Many-body-localized systems fail to thermally equilibrate, so their long-time states are not captured by conventional equilibrium statistical mechanics. At the same time, the existence of interactions allows for highly nontrivial statistical mechanics in these localized systems (11). Indeed, many-body-localized systems can exhibit a phenomenology that runs counter to theorems of equilibrium statistical mechanics, such as ordered phases and phase transitions in one-dimensional systems at energies that correspond to high temperatures (11, 19, 22, 23, 42). In this review, we provide a pedagogical introduction to key features and implications of MBL as well as to the open questions in the field. In fact, this review is almost entirely about theoretical questions. MBL as a subject of experimental research is still in its infancy, so a review of that topic would be premature, and we have not attempted it.

2. CLOSED-SYSTEM MANY-BODY QUANTUM MECHANICS

We are interested in the properties of closed many-body quantum systems with a short-range Hamiltonian H . Unless otherwise specified, we consider a time-independent Hamiltonian, but below we do occasionally consider the Floquet case of a Hamiltonian that is periodic in time. By closed, we mean that the system is not connected to any external environment or to any measuring apparatus. Any environment or measuring apparatus that is coupled to the system should instead be treated quantum mechanically and included as part of the system. We assume, unless otherwise specified, that the Hamiltonian H is local in real space, meaning all interactions are short range.

Unlike many treatments of quantum condensed matter, we do not focus on the ground state and low-lying excited states but instead on highly excited states, which have a nonzero energy density relative to the ground state even in the thermodynamic limit. We also do not restrict ourselves to pure states because the closed system may in the past have been coupled to some other degrees of freedom with which it is still entangled. Arbitrary states of a closed quantum system can be treated using the formalism of probability operators (a.k.a. density matrices) (1). We work in the Schrödinger representation, where the probability operator $\rho(t)$ evolves in time according to

$$\rho(t) = e^{-\frac{iHt}{\hbar}} \rho(0) e^{\frac{iHt}{\hbar}}; \quad i\hbar \frac{d\rho}{dt} = [H, \rho]; \quad \text{Tr}\{\rho\} = 1. \quad 1.$$

Meanwhile, all other operators \hat{O} are time-independent, and the expectation value at time t of an observable O with corresponding operator \hat{O} is $\langle \hat{O} \rangle_t = \text{Tr}\{\hat{O}\rho(t)\}$.

For specificity, we now consider a quantum system of N two-state systems. These two-state systems need not be spin-1/2s, but it is convenient to refer to them as spins. Each spin is located at a point in real space; we consider one, two, or three spatial dimensions. The spins may be randomly located, as in Reference 10, or arrayed on a regular lattice. The locality that localization refers to is in real space. Each spin has a local space of pure states consisting of two states and all of their complex linear combinations. Each spin also has four linearly independent operators that can act on it. These can be represented as 2×2 matrices: the identity matrix \mathbb{I}_i

and the three Pauli matrices σ_i^x , σ_i^y , and σ_i^z for spin i . A general mixed state ρ_i of this spin is a linear combination of these operators. More generally, one could consider a local spin with q orthonormal pure states and q^2 different operators (including its identity operator).

The many-body pure-state space of the full system of N spins is the outer product of the pure-state spaces of each spin (**Figure 1**). A convenient (and conventional) basis set for this space is the 2^N simultaneous eigenstates of all of the $\{\sigma_i^z\}$. This basis set has a geometry: It can be represented as the 2^N corners of an N -dimensional hypercube. Nearest-neighbor points on this hypercube are states that differ by flipping only one spin. Each axis of the hypercube represents one spin, so in general this hypercube does not have any rotational symmetry because each of its axes represents a different point in real space. The Hamiltonian of the many-body system can be represented as a single quantum particle hopping on this hypercube, with only short-range hopping. But a many-body state that is localized need not appear localized on this hypercube: For example, a many-spin product pure state that is localized so each spin points in its own particular direction on its own Bloch sphere, with the directions selected randomly and uniformly on the Bloch spheres, appears quite extended when represented on this hypercube. Such a state only looks localized on the hypercube if the basis choice for each spin corresponds to the state in which that spin is localized. Note that the geometry of this hypercube is not the geometry of the space (the Hilbert space) of the full system's pure states. The latter space is a 2^N -dimensional complex vector space, and each of its 2^N axes is one of the corners of this hypercube. Thus, our system has at least three spaces that can be considered: real space, the hypercube of many-body basis states, and the Hilbert space.

A convenient complete basis for the full system's operators is given by outer-product spin operators of the form, e.g.,

$$\sigma_1^x \otimes \sigma_2^z \otimes \mathbb{I}_3 \otimes \dots \otimes \sigma_N^y, \quad 2.$$

where every spin contributes to the product either its identity operator or one of its Pauli operators. Thus, the system has 4^N linearly independent operators that can operate on its states and from which one may make its mixed states. Again, if we instead have a local spin i with $q_i > 2$ pure states, it contributes a factor of q_i^2 to the number of operators.

Introducing some terminology: A k -local operator is an operator of the above form in which k of the entries are not identity operators (i.e., a product operator that acts nontrivially on only k spins). Such an operator can, for example, be a hopping of range k on the hypercube of many-body basis states. Any linear combination of k -local operators is also termed a k -local operator. A global operator is a k -local operator in which k is of order N . In contrast, an operator that is local in real space (henceforth referred to simply as a local operator) is an operator in which k is of order one, and the nonidentity Pauli operators act only on spins that are all within distance of order one of each other.




Figure 1

(a) The system of interest can be represented as a set of spins on a lattice. (b) Each spin has a two-dimensional-state space, which can be represented on a Bloch sphere. The many-body, pure-state space for the full system consists of the outer product of the pure-state spaces of each spin, as illustrated in panel *b*.

The Hamiltonian H of our system is a sum of local operators. The system may have a few other extensive conserved quantities that are also sums of local operators. Examples include spin and particle number. Like the energy, these quantities can be transported by the system's dynamics. These extensive conserved operators commute with H in a way that is not fine-tuned to the details of H . We do not consider traditional translationally invariant integrable systems that have an infinite sequence of extensive conserved quantities that are sums of local operators. Such integrable systems are special cases and are presumably not robust to arbitrary small local changes in H . We note that the generalized thermalization of such integrable systems to the generalized Gibbs ensemble is an interesting and well-studied subject (46–48) that we do not review here.

Any closed many-body quantum system has many conserved quantities that are given by global operators. A complete linearly independent set of operators that commute with H consists of operators projecting on to the exact many-body eigenstates of H . These projection operators are all global operators, but the Hamiltonian itself is a linear combination of these projection operators and is, in contrast, a sum of local operators (as is also, trivially, the identity operator). There is a certain sense in which local operators are physical and observable, whereas global operators are not (49): Another degree of freedom, such as a measuring apparatus, can realistically couple to a quantity that is represented by a local (or k -local) operator, whereas in the thermodynamic limit this is not feasible for a quantity that is represented by a global operator. Note also that in the thermodynamic limit almost all of this system's full set of 4^N operators are global and thus are unphysical in this sense.

3. QUANTUM THERMALIZATION

3.1. What Is Thermalization?

We are now ready to discuss the apparent paradox of quantum thermalization. A quantum system in thermal equilibrium is fully characterized by a small number of parameters (temperature, chemical potential, etc.: one parameter for each extensive conserved quantity), suggesting that the process of going to thermal equilibrium is associated with the erasure of the system's memory of all other details about its initial state. However, unitary time evolution cannot erase information, and thus all quantum information about the initial state must be preserved within the (closed) system for all times.

The resolution to this apparent paradox is that the memory of the local properties of the system's initial state is not erased by unitary time evolution but is instead hidden if the system thermalizes. Spreading of quantum entanglement moves the information about the initial state so that at long time it is inaccessible because recovering that information would require measuring global operators. This is the process of decoherence. In particular, if we consider only a specific subsystem that is a small fraction of the full system, then thermalization means that at long times the state of this subsystem is as if it were in thermal equilibrium in contact with a reservoir characterized by a temperature T , a chemical potential μ , etc., because in fact it is, with the reservoir being the remainder of our closed system (Figure 2). It is this ability of quantum systems to act as reservoirs for their subsystems that underpins equilibrium quantum statistical mechanics.

We now provide a somewhat more precise description of quantum thermalization. To keep the discussion as simple as possible, we consider a closed system that does not have any extensive conserved quantities other than energy, so that if it thermalizes, the thermal state is described by one parameter, the temperature. A generalization to systems with a few more conserved quantities does not need substantial additional concepts that are not present in this simpler case. The




Figure 2

(a) Conventional quantum statistical mechanics assume that the system of interest is coupled to a reservoir (or bath) with which it can exchange energy and particles. (b) Here, we are interested in the statistical mechanics of a closed quantum system undergoing unitary time evolution. There is no external reservoir. (c) It can be useful to partition the closed quantum system into a (A) subsystem and (B) everything else. If the system quantum thermalizes, then the region (B) is able to act as a bath for the subsystem (A).

interactions in the system’s Hamiltonian must connect all of its degrees of freedom so that the system does not contain any subsystems that are themselves isolated closed subsystems not in contact with the remainder of the system. We partition the full quantum system into a subsystem, A, and its environment, B, which contains all the degrees of freedom not in A. We need to take the thermodynamic limit on the number of degrees of freedom in B such that in this limit the fraction of the full system’s degrees of freedom that are in A goes to zero. Any choice of subsystem A is acceptable, as long as the degrees of freedom within A are defined by k -local operators with finite k . A concrete case can be that A is a fixed compact subregion (in real space). However, different choices, such as a subsystem in momentum space, a set of degrees of freedom that are well separated in real space, or even a single degree of freedom, also constitute acceptable subsystems. For clarity, we consider a case in which A is a fixed compact subregion of the full system and B is taken to infinite volume by adding more degrees of freedom that are in the limit arbitrarily far from A. Thus, as we take the thermodynamic limit on B, we have a sequence of systems and their Hamiltonians, with the number of degrees of freedom increasing without limit. In the concrete case mentioned, the changes to H as the system’s size is increased are only at the locations far from A where the new degrees of freedom are being added to B.

For each system in this sequence, as we take the thermodynamic limit we need to consider a set of initial states, $\rho(t=0)$. Let us consider initial states that, if they do thermalize, thermalize to a given temperature T . Each system in our sequence has an equilibrium expectation value of the total energy $\langle H \rangle_T$ at temperature T . We consider a sequence of initial states such that the mean-square deviation within each initial state of the total energy of the system from $\langle H \rangle_T$ grows no faster than the volume of the full system. Although we thus fix the average energy density to be at its equilibrium value, we do not otherwise constrain how this energy is initially distributed, and we are particularly interested in initial states in which that initial distribution of energy is far from equilibrium, because these are a type of initial state that may fail to thermalize. Each system and each initial state is then time-evolved, as found in Reference 1. The probability operator $\rho_A(t)$ (a.k.a. reduced density matrix) of subsystem A at time t is obtained from $\rho(t)$ of the full system by taking a partial trace over all of the degrees of freedom in B: $\rho_A(t) = \text{Tr}_B\{\rho(t)\}$. The same system at equilibrium at temperature T has Boltzmann probability operator $\rho^{(eq)}(T) = Z^{-1}(T)\exp(-H/k_B T)$ for the full system and thus $\rho_A^{(eq)}(T) = \text{Tr}_B\{\rho^{(eq)}(T)\}$ for the subsystem. The system thermalizes for this temperature if in the long-time and large-system limit $\rho_A(t) = \rho_A^{(eq)}(T)$ for all subsystems A (50–53). These two limits must be taken together: For a finite system, the dynamics are quasiperiodic, so $\rho(t)$ does not have a long-time limit, whereas for finite time, the diffusive transport in a thermalizing system only reaches a finite distance.

Thermalization is of particular interest for initial states that are well out of equilibrium. These are atypical initial states because at equilibrium the system is in typical states. Thus, when we say a system thermalizes, for this to be an interesting statement it must certainly apply to some atypical initial states. The conventional story in equilibrium statistical mechanics is that all out-of-equilibrium initial states at this energy thermalize in the limit of infinite time to equilibrium at the corresponding temperature. Thus, it seems reasonable to assume that if a system does thermalize for a given temperature, this means all initial states at that energy thermalize. This is a strong statement with some strong consequences. It seems to be well beyond what can be proved for any generic system, but it seems to be consistent with what we know so far about systems that do thermalize.

It is interesting to consider a many-body Floquet system, a system with a time-dependent Hamiltonian that is periodic in time with period τ , so $H(t) = H(t + \tau)$. In such systems, energy is only conserved modulo $2\pi\hbar/\tau$, so there is no conserved energy density. Let us consider a Floquet system that has no conserved densities at all and thus nothing to transport. A periodically driven system with an unbounded Hilbert space cannot really thermalize; instead, it can absorb energy from the periodic drive without limit. However, when a periodically driven system with a bounded Hilbert space thermalizes, it thermalizes to infinite temperature (45), at which point all pure states of any small subsystem are equiprobable. When such a system thermalizes, it still serves as a reservoir for its subsystems, but a reservoir for what, given that there is no conserved quantity that is being exchanged between the subsystem and the reservoir. This suggests that the essential function of a reservoir is not as a source and sink of energy, particles, or other such conserved quantities. Instead, the most basic function of the reservoir may be to provide other quantum degrees of freedom that the subsystem gets so entangled with that no information about the initial state of the subsystem remains locally observable.

3.2. The Eigenstate Thermalization Hypothesis

If a system at a given temperature does indeed thermalize for every such initial state $\rho(t=0)$, then it is quite instructive to consider initializing the system in a pure state that is one of the many-body eigenstates of H . The time evolution of the system then becomes trivial: $\rho(t) = \rho(0)$, so thermalization of all initial states implies that all many-body eigenstates of H are thermal. This statement is known as the eigenstate thermalization hypothesis (ETH) (50–53). Before we discuss the ETH in more detail, we must emphasize that we are not considering the exact eigenstates of H because these are realistic states of a many-body system. On the contrary, they are impossible to prepare in the laboratory as initial states. The initial states of a many-body system that can actually be prepared are not even pure states, much less eigenstates. As always, the focus on eigenstates is instead because they are an essential tool in understanding the dynamics. When written as a density matrix using the eigenbasis $\{|n\rangle\}$ of H , the dynamics of $\rho(t)$ is simple: Its diagonal terms $\rho_{nn} = \langle n|\rho|n\rangle$ are constant, whereas its off-diagonal terms $\rho_{nm} = \langle n|\rho|m\rangle$ each simply precess in the complex plane at a constant rate given by the difference in energy between the two eigenstates involved: $\rho_{nm}(t) = \rho_{nm}(0) \exp[i(E_m - E_n)t/\hbar]$.

To define the ETH a little more precisely, consider an eigenstate $H|n\rangle = E_n|n\rangle$. Its energy E_n is the thermal equilibrium energy at temperature T_n , so $E_n = \langle H \rangle_{T_n}$. If the full system is in this eigenstate, then $\rho = \rho^{(n)} = |n\rangle\langle n|$, and thus $\rho_A^{(n)} = \text{Tr}_B\{|n\rangle\langle n|\}$ is the state of subsystem A. The ETH asserts that in the thermodynamic limit the subsystem is at thermal equilibrium: $\rho_A^{(n)} = \rho_A^{(eq)}(T_n)$. One noteworthy consequence of this is that the entropy of entanglement, $S_{AB} = -k_B \text{Tr}_A \left\{ \rho_A^{(n)} \log \rho_A^{(n)} \right\}$, between A and B in this eigenstate of the full system is equal to the equilibrium thermal entropy of the (smaller) subsystem A. For eigenstates with temperature $T_n \neq 0$, this entropy is proportional to the volume of A. Thus, the entanglement entropy in thermal eigenstates obeys volume-law scaling, as it does also in any thermal pure state of the full system.

Another requirement of the ETH is that the matrix elements of subsystem operators between distinct eigenstates vanish strongly enough in the thermodynamic limit (53, 54). This is needed to ensure that temporal fluctuations of $\rho_A(t)$ vanish [as opposed to the weaker scenario in which only the time average of $\rho_A(t)$ is thermal].

For many-body Floquet systems, the dynamically stable many-body states of the full system that play the role of the Hamiltonian's eigenstates are the eigenstates of $U(\tau)$, the unitary operator that takes the full system forward in time by one period τ . Of course, these eigenstates of $U(\tau)$ are not completely stationary; they only return to the same state after a full period of the drive. For a system with no conserved densities, the equilibrium state of a subsystem has $\rho_A^{(eq)}$ be simply equal to the identity operator on that subsystem (times a normalizing factor), so all pure states of the subsystem are equally probable. Thus, when the ETH applies to such Floquet systems, it means that in the thermodynamic limit all eigenstates of the full system give equal probabilities for all possible pure states of any finite subsystem.

The ETH is a hypothesis. It is not true for one broad class of systems, namely those that are many-body Anderson localized, as we discuss below. For systems in which the ETH appears to be true, this is difficult to thoroughly test numerically because it requires obtaining the many-body eigenstates of the system's Hamiltonian from exact diagonalization and extrapolating to the thermodynamic limit. In order to have thermalization of all the initial states that can realistically be prepared, it seems that one does not need the ETH to be true for absolutely all eigenstates; a weaker almost all should suffice. And if we look at the numerical results, they do strongly support the proposal that there are systems in which at least almost all eigenstates obey the ETH (17, 53–65). But the simpler and, it seems, more plausible scenario is that if the ETH is true for a given system at a given temperature, then it is true for all eigenstates. If there were certain rare eigenstates that violate the ETH even though they are essentially degenerate with the typical eigenstates that obey the ETH, we would need to understand what is so special about these rare eigenstates.

We note that if the ETH is true and the state ρ of the system is diagonal in the energy eigenbasis, then all subsystems are at thermal equilibrium. This prompts the question: When the ETH is true, how does one construct out-of-equilibrium states? The answer is that out-of-equilibrium states have special structure off the diagonal when ρ is written in the energy eigenbasis. Thus, they have special coherence patterns between eigenstates of different energies. Quantum thermalization of an out-of-equilibrium initial state requires that the contributions of this off-diagonal coherence to local observables must vanish at long times. This happens due to dephasing: Although the diagonal terms in ρ (in the energy eigenbasis) are time-independent, the off-diagonal terms have phase oscillations at frequencies set by the energy differences between the corresponding eigenstates. Thus, unitary time evolution scrambles the phases of the off-diagonal terms in ρ such that at long times and in the thermodynamic limit their contributions to any local observables come with effectively random phases and thus cancel. In this sense, we see that in the energy eigenbasis equilibration is simply dephasing.

The ETH motivates the introduction of a new set of ensembles to use in quantum statistical mechanics, namely the single-eigenstate ensembles that each consist of a single eigenstate of the full system's Hamiltonian. When the ETH is true, these ensembles all give the correct thermal equilibrium properties of subsystems, just like the traditional statistical mechanical ensembles. The single-eigenstate ensembles may be viewed as the limiting case of the microcanonical ensemble in which the energy window has been reduced to the limit at which it contains only one eigenstate. The full payoff of introducing these new ensembles and the resulting single-eigenstate statistical mechanics becomes clear when one considers systems that do not obey the ETH, as we discuss below.

Although the ETH and quantum thermalization appear to apply to a large class of closed quantum systems, not all systems quantum thermalize. One well-known exception is traditional

integrable systems, which possess an infinite set of extensive conserved quantities. It has, however, been argued that such integrable systems exhibit their own version of quantum thermalization to a generalized Gibbs ensemble (GGE), and they have their own version of the ETH (46–48). The focus of the remainder of this review, however, is on a class of systems that fail to quantum thermalize in any sense and on where the many-body eigenstates violate the ETH. These are Anderson-localized systems.

4. LOCALIZED SYSTEMS

The concept of localization was first introduced by Anderson (10) and applies to systems with quenched disorder (see, however, Section 6.1 for a discussion of the possibility of MBL in lattice systems with translationally invariant Hamiltonians). The label localization is used for at least three different situations: Most experimental work on localization to date is about regimes near ground-state quantum phase transitions (e.g., metal-insulator transitions). This is not a topic of this review. Most theory work on localization to date is concerned with noninteracting particles in a random potential (or waves in random media without nonlinearities). We briefly review this work in Section 4.1, but it is also not a main topic of this review. The focus of this review is on interacting many-body systems, and we are interested not in ground states or the low-energy limit but instead in highly excited states of such systems at energies that would correspond to nonzero temperature if the system thermalized. The discussion of MBL begins in Section 4.2 and occupies the remainder of this review.

4.1. Single-Particle Localization

We briefly review single-particle localization (for a more complete discussion, see, e.g., Reference 66). The essential physics of single-particle localization can be illustrated with a tight-binding model of a single quantum particle hopping on an infinite lattice, with Hamiltonian

$$H = t \sum_{\langle ij \rangle} (c_i^\dagger c_j + c_j^\dagger c_i) + \sum_i U_i c_i^\dagger c_i, \quad 3.$$

where U_i is a static random onsite potential, $t \neq 0$ is a nearest-neighbor hopping, and c_i^\dagger creates a particle on the site i . Consider the motion of a single particle in this system. In three or more dimensions, and for weak-enough disorder, the eigenstates of this Hamiltonian can be extended with weight on all sites, and diffusive dynamics of a particle can be initialized in a wave packet composed of such extended eigenstates. However, in one or two dimensions (and in three or more dimensions with strong-enough disorder), the eigenstates are all exponentially localized, with wavefunctions that have the asymptotic long distance form $\psi_\alpha(\mathbf{r}) \sim \exp\left(-\frac{|\mathbf{r} - \mathbf{R}_\alpha|}{\xi}\right)$, where ξ is the localization length, which depends on the disorder strength and the energy. This state α is localized near position \mathbf{R}_α , and a particle remains localized near the location at which it is initially introduced. In three or more dimensions, the transition between localized and extended states happens via special critical states at the mobility edge, which display power-law localization.

A straightforward basis transformation recasts this Hamiltonian into the simple form

$$H = \sum_\alpha E_\alpha c_\alpha^\dagger c_\alpha, \quad 4.$$

where c_α^\dagger creates a particle in the single-particle eigenstate $|\alpha\rangle$, and E_α is the eigenenergy of a particle occupying this state. We have written this Hamiltonian in second-quantized form, so it is also

a many-particle Hamiltonian, although still without interactions between the particles. Its many-particle eigenstates are simple product states in terms of the single-particle eigenstates, and can be labeled by the occupation numbers of all the various single-particle eigenstates $|\alpha\rangle$. For systems in which at least some of the single-particle eigenstates are localized, almost all of these many-particle eigenstates violate the ETH. As an example to show the lack of quantum thermalization, the system is initialized with a spatially nonuniform density of particles in the localized states over a large length scale: This nonuniform initial density pattern then survives for all times because those particles are localized.

Although particles hopping in a random potential provide the best-known example of single-particle localization, the phenomenon is more general and occurs also in systems of spins. For example, a spin-1/2 version of single-particle localization arises in a system governed by the spin Hamiltonian

$$H = \sum_i b_i \sigma_i^z + \sum_{ij} J_{ij} \sigma_i \cdot \sigma_j. \quad 5.$$

The onsite magnetic fields b_i are static random variables, e.g., taken from a continuous probability distribution of width $W > 0$, and the spin hoppings and interactions J_{ij} are strictly short range in real space (for specificity, take $J_{ij} = J \neq 0$ for nearest neighbors and zero otherwise). One then considers an initial condition in which there is a single spin up (e.g., on site i) with all other spins down and asks about the dynamics of this up spin.

If there are any localized single-particle states with weight at site i , there is a nonzero probability that the site i still hosts an up spin even at infinite time, i.e., memory of the initial conditions is preserved in a local observable for infinite times. This manifestly constitutes a failure of quantum thermalization and is quite similar to the context in which single-particle localization was first established in Reference 10.

4.2. Many-Body Localization

We now turn to MBL: localization with interactions. This physics can be studied in models with mobile particles. Because it occurs at high energies, it appears that there are no major differences between MBL of bosons and fermions. But the physics of the MBL phenomenon is most simply exposed in the context of spin models. The model above (Equation 5) is fine to illustrate MBL; we simply consider states in which the densities of up and down spins are nonzero, instead of considering only one flipped spin. We are interested in whether or not a system governed by the above Hamiltonian (Equation 5) quantum thermalizes for arbitrary initial conditions. We therefore ask whether the many-body eigenstates of the above Hamiltonian obey the ETH, given that this is a necessary condition for quantum thermalization.

At $J = 0$, the many-body eigenstates of Equation 5 are simply product states of the form $|\sigma_1^z\rangle \otimes |\sigma_2^z\rangle \otimes \dots$, and the system is fully localized. For nonzero J , in the regime $J \ll W$ one can construct the many-body eigenstates perturbatively in small J (12–15). Because in this regime the typical level splittings between nearest-neighbor sites are much larger than the interactions J , the states on different sites are typically only weakly hybridized. This line of argument (similar to that employed by Anderson in Reference 10) leads one to conclude that for sufficiently strong disorder $W \gg J$, DC spin transport and energy transport is absent, and quantum thermalization therefore does not occur at any order in perturbation theory (15). Although this argument is perturbative and is limited to the weak interaction regime, extensive numerical evidence [mostly for one-dimensional systems (16, 17, 19)] suggests that while the high-entropy eigenstates of Equation 5 do obey the ETH for weak disorder (small W), for strong-enough disorder all of the eigenstates violate ETH.

Moreover, the violation of ETH apparently occurs even for strong interactions outside the regime in which the perturbation theory can be controlled.

There is a quantum phase transition as one varies the disorder strength or the energy density between the thermal phase, in which the system quantum thermalizes and we expect all the eigenstates to obey the ETH, and the many-body-localized phase, wherein all the eigenstates do not obey the ETH and some memory of the local initial conditions can survive in local observables for arbitrarily long times. Many questions about the nature of this phase transition remain open (see Section 6.2). This transition is an eigenstate phase transition, marked by a sharp change in properties of the many-body eigenstates and thus in the dynamics of the system, so this transition is visible if one studies the system using the single-eigenstate ensembles. However, this transition is invisible to equilibrium thermodynamics and to the traditional statistical mechanical ensembles because they average over many eigenstates. Indeed, the MBL transition marks the breakdown of the applicability of equilibrium quantum statistical mechanics to the system's long-time properties.

Although localization is usually discussed for systems with static randomness, it has long been known that nonrandom systems with quasiperiodicity can support single-particle localization. This has recently been demonstrated to also remain true for MBL (67).

Although the original idea of localization came from considering spin systems (10), and spin models provide simple examples in which to explore MBL, important, more-recent work about MBL also considered systems of fermions: In Reference 13, it was pointed out that a system of interacting fermions in zero dimensions (a quantum dot) can be approximated by a single-particle localization problem on a Cayley tree. This result was exploited in References 14 and 15 to show that single-particle localization in spatial dimensions $d \geq 1$ is robust to weak nonzero interactions, to all orders in perturbation theory. A feature that can occur is a many-body mobility edge, which was emphasized in Reference 15, at an extensive energy in systems in which some but not all many-body eigenstates are localized. The usual behavior in the energy regime that corresponds to positive temperature is that all many-body eigenstates with energy density above the mobility edge (thus in the thermal phase) obey the ETH, whereas all eigenstates with energy density below the mobility edge are in the localized phase. Models with an inverted mobility edge in which localization occurs only for those eigenstates above a critical energy density can also be constructed by making use of a model in which the single-particle level spacing increases with energy (41). We are not aware of any models exhibiting multiple mobility edges in the positive temperature range, but there does not seem to be any obvious reason why such models cannot be constructed. The behavior of MBL systems with a mobility edge is more complex than that of systems in which all eigenstates are localized. In particular, one may worry about rare regions within a localized state that have a local energy density close to that of the mobility edge (a new type of quantum Griffiths phenomenon).

We close this section by emphasizing one key feature of the MBL phenomenon: Closed MBL systems do not quantum thermalize, so some memory of the local initial conditions is preserved in local observables for arbitrarily long times. This implies that the DC conductivity of any conserved densities must be strictly zero in the MBL phase. Although the vanishing of DC transport is a useful diagnostic for MBL in certain systems, it is not the key distinguishing feature. Indeed, one can consider a many-body Floquet system with no conserved local densities, meaning it has no meaningful DC transport properties. However, such Floquet systems may have a quantum phase transition between thermal and MBL phases (see, e.g., 44, 45).

4.3. A Phenomenology of Many-Body-Localized Systems

We begin this section by tabulating some properties of the thermal (nonlocalized), single-particle localized, and many-body-localized phases (Table 1). The first three lines of Table 1 follow from

Table 1 A list of some properties of the many-body-localized phase, contrasted with properties of the thermal and the single-particle-localized phases

Thermal phase	Single-particle localized	Many-body localized
Memory of initial conditions hidden in global operators at long times	Some memory of local initial conditions preserved in local observables at long times	Some memory of local initial conditions preserved in local observables at long times
Eigenstate thermalization hypothesis (ETH) true	ETH false	ETH false
May have nonzero DC conductivity	Zero DC conductivity	Zero DC conductivity
Continuous local spectrum	Discrete local spectrum	Discrete local spectrum
Eigenstates with volume-law entanglement	Eigenstates with area-law entanglement	Eigenstates with area-law entanglement
Power-law spreading of entanglement from nonentangled initial condition	No spreading of entanglement	Logarithmic spreading of entanglement from nonentangled initial condition
Dephasing and dissipation	No dephasing, no dissipation	Dephasing but no dissipation

the discussion in Section 4.2. However, explaining the rest of this table requires a little more discussion, which we now provide.

Let us assume for specificity that we have a system of N -local, two-state degrees of freedom $\{\sigma_i\}$, which we refer to as the p-bits (p, physical). These could be the spins from Reference 9 or, e.g., the occupation numbers of localized single-particle orbitals in a system of fermions in a random potential. An analogous argument can be constructed for objects with more than two states, but we stick to this two-state example for specificity. Let us further assume that the p-bits are governed by a Hamiltonian with quenched disorder and strictly short-range interactions. For strong-enough disorder, such a Hamiltonian can be in the fully many-body-localized (FMBL) regime, wherein all the many-body eigenstates of the Hamiltonian are localized. It was argued in References 26–29 that in this FMBL regime, one can define a set of localized two-state degrees of freedom, with Pauli operators $\{\tau_i\}$, henceforth called l-bits (l, localized) such that the Hamiltonian when written in terms of these new variables takes the form

$$H = E_0 + \sum_i \tau_i^z + \sum_{ij} J_{ij} \tau_i^z \tau_j^z + \sum_{n=1}^{\infty} \sum_{i,j,\{k\}} K_{i\{k\}j}^{(n)} \tau_i^z \tau_{k_1}^z \dots \tau_{k_n}^z \tau_j^z, \quad 6.$$

where the sums are restricted so that each term appears only once, and E_0 is some constant energy offset that may be zero and that has no relevance for the closed system's dynamics. The typical magnitudes of the interactions J_{ij} and $K_{i\{k\}j}^{(n)}$ fall off exponentially with distance, as do their probabilities of being large.

The intuition underlying this l-bit Hamiltonian (Equation 6) is that in the localized phase, because there is no transport, there should be a set of localized conserved charges that are constants of motion of the system. For example, for noninteracting particles moving in a disordered potential, these constants of motion would be the occupation numbers of the localized single-particle orbitals (Figure 3). For weakly interacting systems, the l-bits τ_i should have substantial overlap with the bare p-bits σ_i , and indeed may be viewed as dressed p-bits, with a dressing that falls off exponentially in real space. The existence of long-range interactions between l-bits follows



Figure 3

An illustration of the dynamical behavior of the l-bits (l, localized). The l-bits are a set of spins whose z component does not change but which precess about the z axis at a rate determined by the effective interactions with all other l-bits. This picture can be used to understand, e.g., the logarithmic spreading of entanglement in the fully many-body-localized phase.

from the fact that although the Hamiltonian only couples p-bits that are nearby in real space, each of those p-bits has nonzero (but typically exponentially small in the distance) weight on distant l-bits.

One appealing approach to constructing l-bits is to start with p-bits and then to make operators that commute with each other and with the Hamiltonian by adding the appropriate dressing order by order in perturbation theory in the p-bit interactions. This approach fails to give a unique result away from the limit of weak interactions; also, for some l-bits, this approach fails to give a unique result even in the weakly interacting regime because of the appearance of resonances (degeneracies in the perturbation theory) that make the definition of the l-bits ambiguous. Nevertheless, it was argued in References 26–29 that localized l-bits do exist for FMBL systems, and each such l-bit has overlap with distant p-bits that are typically exponentially small in the distance. Something essentially equivalent to this is proven in Reference 18. No such construction is possible in the thermal phase. Whether any such construction is possible in the localized phase for MBL systems with a many-body mobility edge remains an open question.

The l-bit Hamiltonian (Equation 6) provides a useful tool for describing various properties of FMBL systems (Figure 3). The eigenstates of the Hamiltonian (Equation 6) are simply the simultaneous eigenstates of all the $\{\tau_i^z\}$. In a generic state, the dynamics of a single l-bit are in a certain sense trivial: Each l-bit precesses around its z axis at a rate set by its interactions with all other $\{\tau_i^z\}$. In a generic state where all the $\{\tau_i^z\}$ are uncertain, this precession of l-bits produces entanglement and dephasing. Nevertheless, dynamically there are no flips of the $\{\tau_i^z\}$ and thus no dissipation. As a result, the dephasing can, in principle, be reversed by spin-echo procedures, so MBL systems can, in principle, be used to store and retrieve quantum information. However, to fully reverse the dephasing of a particular l-bit, we need to be able to manipulate only that l-bit, whereas in a general experiment one can only access the p-bits. Recent work (68) suggests that high-fidelity spin-echo measurements are indeed possible even when one has access to only p-bits, at least in the perturbative regime where the typical l-bits are weakly dressed p-bits. How strong a spin-echo signal can be obtained by doing the echo procedure on a bare p-bit away from the perturbative regime remains an open question.

The structure of the Hamiltonian (Equation 6) can also be used to understand the logarithmic spreading of entanglement in the FMBL phase (26, 27), observed numerically in References 30 and 69. It is useful to first consider how entanglement spreading works in thermalizing and in single-particle localized systems and to contrast these with many-body-localized systems: In thermalizing systems, the interaction of two p-bits, A and B, generically causes them to become entangled with each other. The subsequent interaction of bits B and C generically produces entanglement not just between B and C but also between A and C. As a result, entanglement spreads ballistically in some systems (70), with a speed akin to a Lieb-Robinson velocity (71). (Unpublished results from R. Vosk, E. Altman, and D.A. Huse suggest that the entanglement spread can instead be a sub-ballistic power law of time in one-dimensional thermalizing systems if weak links due to rare, almost-insulating regions are common enough to render the entanglement spread sub-ballistic.) In

contrast, in single-particle localized systems, there are no interactions between l-bits, so the dynamics do not generate any entanglement between l-bits.

The entanglement spreading in FMBL systems is intermediate between thermalizing and noninteracting localized systems. Entanglement spreading does occur because the system is interacting. However, two l-bits can get entangled only by their direct interaction and not through a mutual interaction with a third l-bit. This is because the interaction of A with B depends only on the τ^z value for l-bit B, and the τ^z value for l-bit B is unaffected by its interaction with a third l-bit C, given that the τ^z value for l-bit B is a constant of motion. In other words, l-bits become entangled only through their direct interaction because the l-bit Hamiltonian (Equation 6) has no dissipation (no spin flips).

To quantify the interaction between l-bits, we define the effective interaction between l-bits i and j as

$$J_{ij}^{eff} = J_{ij} + \sum_{n=1}^{\infty} \sum_{\{k\}} K_{i\{k\}j}^{(n)} \tau_{k_1}^z \dots \tau_{k_n}^z. \quad 7.$$

Note that this interaction depends on all the other $\{\tau_k^z\}$. In the localized phase, we expect this interaction to typically fall off with the distance L between the two l-bits, as $J^{eff} \sim J_0 \exp(-L/\xi)$. If these two l-bits are initially not entangled, this interaction entangles them after a time t such that $J^{eff} t \geq 1$. As a result, after a time t , the system's dynamics produces entanglement between all l-bits within a distance $L \sim \xi \ln(J_0 t)$ of each other. This is the origin of the logarithmic growth of entanglement with time within the MBL phase.

Note that there are exponentially many multispin terms contributing to the effective interaction (Equation 7), and only a single two-spin term. Thus, for long distances, the effective interaction is dominated by the multispin terms, and the individual multispin terms in the sum (Equation 7) typically fall off exponentially with a decay length much less than ξ . The dominant multispin terms have all the k_i near the straight line between sites i and j . Flipping a single spin between i and j changes the sign of an appreciable fraction of the multispin terms such that the effective interactions change dramatically from one many-body eigenstate to the next—a form of temperature chaos analogous to spin glasses (72). The effective interaction relevant for a particular experiment is obtained by taking the appropriate average of Equation 7 over the particular state in which the system is prepared: This is discussed more fully in Reference 29. The effective interaction can, in principle, be measured experimentally through a double electron-electron resonance-type protocol, as discussed in Reference 68.

The l-bit phenomenology introduced above works for systems that are FMBL, i.e., all of the many-body eigenstates of the Hamiltonian are localized. Whether an analogous construction exists for Hamiltonians that have both extended and localized eigenstates, separated in energy by a many-body mobility edge, remains an open question. Naïvely, one might think that a similar construction, restricted to localized states with energies below the mobility edge, might have a chance of success. However, the proper treatment of rare regions in which the local energy density approaches the many-body mobility edge (a new type of Griffiths phenomenon) complicates this line of reasoning such that the development of an analogous phenomenology for systems with a many-body mobility edge remains an open problem.

4.4. Open Systems, Local Spectra

Although perfectly isolated MBL systems are a useful theoretical idealization, any actual physical system always has some small coupling to its external environment (a bath). We therefore are

interested in not just the limit of perfect isolation but also how this limit is approached and what aspects of MBL phenomena survive in the presence of imperfect isolation from the bath (Figure 4). Provided that the bath is large enough and itself thermalizes, the exact eigenstates of the coupled system and bath for nonzero coupling generically obey the ETH. However, localization of the system is recovered in the limit when the coupling to the bath is taken to zero. This situation was explored in Reference 73, which considered an MBL system coupled to a bath.

Some quite instructive probes of localization behavior are the frequency-dependent spectral functions $A(\omega)$ of local operators of the system (73). If the system is localized and isolated, and (unrealistically) is in an exact many-body eigenstate of the system's Hamiltonian, then the local spectra consist of discrete delta functions, each of which arises from flipping some set of the local l-bits near the location being probed. However, the actual state of a real physical system is not pure; instead it is a mixed state that populates a set of eigenstates with extensive entropy. If the system is at equilibrium with a thermal bath, then this mixed state is the Boltzmann-Gibbs distribution. In these latter cases, for an interacting MBL system the number of delta functions in a local spectral function of the isolated system is exponentially large in the system's volume because the spectral lines change in frequency between eigenstates owing to the interactions between l-bits. The most robust spectral signature of MBL is a soft gap at zero frequency in the spectra of local operators due to level repulsion between the local excitations (73, 74). The local spectral function of a specific isolated system at a specific location may also show gaps that are peculiar to that location, gaps that are thus not present in spatially averaged spectral functions. In isolated one-dimensional systems with short-enough localization length (referred to in Reference 73 as systems exhibiting strong MBL), there are spectral gaps at an infinite hierarchy of scales, with the local spectrum being discrete and having a structure analogous to a Cantor set. In more than one dimension, or in one-dimensional systems with longer localization length (referred to in Reference 73 as systems exhibiting weak MBL), the discrete lines in the local spectrum of a mixed state fill in to make a piecewise continuous spectrum, which has only a finite number of gaps (including the soft gap at zero frequency that is always present in isolated MBL systems).

Introducing a nonzero coupling to a thermal bath causes the spectral lines to broaden into Lorentzians, erasing all structure in the local spectral functions on energy scales smaller than the Lorentzian width. The Lorentzian width goes continuously to zero in the limit of perfect isolation, revealing the above-discussed spectral structure. Meanwhile, as the coupling to the bath is increased, the Lorentzian width also increases, erasing structure on ever larger energy scales. When all structure due to MBL in the local spectral functions is erased by coupling to a bath, the system ceases to show any signs of being many-body localized. This general picture can be used to




Figure 4

One can consider an almost perfectly isolated quantum system, i.e., a quantum system that is coupled to an external bath with a weak but nonzero coupling g that couples (weakly) to all of the degrees of freedom in the system.

consider a wide variety of properties of imperfectly isolated MBL systems, including transport behavior and properties as a quantum memory.

These ideas were investigated numerically in Reference 74, which used exact diagonalization to study the behavior of a spin chain (the system) with a Hamiltonian (Equation 5) coupled to a thermalizing spin chain whose eigenstates obey ETH (the bath) with a tunable coupling g . For numerical reasons, the study was restricted to considering eight spins or fewer in the system and eight spins or fewer in the bath. When the system and bath were decoupled, the local spectrum of a spin-flip operator in the system was found to be very different in the thermal and localized phases (Figure 5*a,b*). On coupling a localized system to the bath, it was found that the eigenstates of the combined system and bath become effectively thermal, with the crossover to thermal eigenstates occurring at a coupling that is exponentially small in the size of the bath. In contrast, the local spectrum was found to remain spiky and inhomogeneous even for system-bath couplings where the combined eigenstates were effectively thermal (Figure 5*c,d*). The behavior of the spectrum is illustrated by Figure 5. Note the existence of a soft gap at zero frequency and also the hierarchy of gaps discussed above, and note too how these gaps are gradually filled in as the coupling to the bath is increased.

The discussion in References 73 and 74 assumed that the local bandwidth of the bath was larger than the characteristic energy scales in the system. The behavior in the opposite limit, when the local bandwidth of the bath is smaller than the characteristic energy scales of the system, was




Figure 5

The spectrum of a local spin-flip operator, σ_i^x , of a system governed by a Hamiltonian (Equation 5) coupled to a thermalizing spin chain (a bath) with a coupling g . The spectra are obtained by exact diagonalization, for a system and bath that consist of eight spins each. The spectra are averaged over spatial position in the system, and also over all many-body eigenstates (i.e., the spectra are evaluated in an infinite temperature Gibbs mixed state). Here, w controls the disorder strength in the system. For small w (low disorder), the system is in its (a) thermalizing phase, and the local spectrum is smooth and continuous. For large w (b–d), the system is in its localized phase if isolated. For the coupled finite-size system and bath, the eigenstates become effectively thermal above a coupling $g \approx 0.15$.

Panel *b* shows the local spectrum in the regime where the eigenstates are nonthermal: The local spectrum is highly inhomogeneous and contains a hierarchy of gaps, including the soft gap at zero frequency. Panels *c* and *d* show the local spectrum at coupling where the eigenstates of the coupled system and bath are thermal. We see that although spectral line broadening does smooth out the local spectrum, the local spectrum retains signatures of proximity to localization. Figure adapted with permission from Reference 74.

investigated in Reference 75. A qualitatively similar behavior was found, but with a line broadening that scales as a power-law function of the bandwidth of the bath.

5. LOCALIZATION-PROTECTED QUANTUM ORDER

If we look at the exact many-body eigenstates of finite systems, we can ask about their properties as we take the thermodynamic limit. So far, we have focused on one question: Are the eigenstates thermal or localized? However, we can also ask about whether these eigenstates exhibit symmetry breaking or topological order. For eigenstates that are thermal, the eigenstates are ordered only when the system is ordered at thermal equilibrium because the eigenstates are each individually at thermal equilibrium. However, in the many-body-localized phase, ordering may be present in individual eigenstates that is absent in thermal equilibrium. This is where the new single-eigenstate statistical mechanics can reveal new phases and phase transitions. We highlight in particular the possibility of spontaneous symmetry breaking even below the equilibrium lower critical dimension and the existence of topological order without a bulk gap. In both cases, an ordered phase that would be destroyed at thermal equilibrium by thermal fluctuations is dynamically protected when those fluctuations are localized and thus static. Examples of localization-protected quantum order have been discussed in References 11, 19, and 22–25.

5.1. Example: Ising Spin Chain

One of the simpler examples to illustrate localization-protected quantum order is a random Ising spin chain, with Hamiltonian

$$H = -\sum_{i=1}^L h_i \sigma_i^x - \sum_{i=1}^{L-1} (J_i \sigma_i^z \sigma_{i+1}^z + \lambda_i \sigma_i^x \sigma_{i+1}^x), \quad 8.$$

where the h_i , J_i , and λ_i are nonnegative and are drawn from some probability distributions (e.g., log normal). This Hamiltonian has an Ising Z_2 symmetry that is implemented by the operator $\hat{P} = \prod_{i=1}^L \sigma_i^x$, which rotates all spins by angle π about their x axes.

When the $\{J_i\}$ are all much larger than the $\{h_i, \lambda_i\}$, the ground state of Equation 8 is ferromagnetically ordered. More precisely, for any finite L , there are two nearly degenerate Schrödinger cat ground states $\approx (|\uparrow\rangle \pm |\downarrow\rangle)/\sqrt{2}$ that are linear combinations of macroscopically different states that are magnetized up ($|\uparrow\rangle$) and down ($|\downarrow\rangle = \hat{P}|\uparrow\rangle$). The energy difference between these two eigenstates is exponentially small in L and is due to the high-order process in which the σ^x terms in Equation 8 act to flip the global magnetization between up and down. If we start this system in the initial magnetized (symmetry-broken) state $|\uparrow\rangle$, which is a linear combination of these two ground states, this state has an exponentially small energy uncertainty and thus takes a time that is exponentially long in L to evolve away from the symmetry-broken $|\uparrow\rangle$ initial state under unitary time evolution. The divergence of this timescale exponentially with system size is diagnostic of spontaneously broken symmetry in the thermodynamic limit. Alternative diagnostics for spontaneously broken symmetry in a finite-size system include the spin-spin correlations $\langle \sigma_i^z \sigma_{i+r}^z \rangle$, which show long-range order in both of the ground states, and the system's susceptibility to magnetic fields along the z axis, which diverges exponentially with system size in the ground states.

Although the ground states of the above system can break symmetry, in the absence of static randomness there is no symmetry breaking at energy densities above that of the ground states. Such an excited state has a nonzero density of domain walls, and in the absence of randomness

these domain walls are delocalized and dynamically wander over the entire chain. The presence of such delocalized domain walls means that the equal-time spin correlations do not exhibit long-range order. Equally, a broken-symmetry initial state (e.g., a magnetization pattern) dynamically relaxes through the motion of the domain walls. The absence of symmetry breaking in the exact eigenstates can be viewed as following from the Landau-Peierls theorem. This theorem states that in one dimension there is no spontaneous symmetry breaking at thermal equilibrium at nonzero temperatures and applies here as long as the system is not many-body localized.

5.2. Symmetry Breaking at Nonzero Energy Density with Randomness

Now let's consider the same system (Equation 8) with strong $\{J_i\}$ so that the ground states remain ferromagnetic but now with sufficiently strong randomness so that the domain walls in its excited states are many-body localized. Let $|\phi\rangle$ be a particular pattern of σ^z magnetization, namely a given pattern of domains and localized domain walls. $\hat{P}|\phi\rangle$ has precisely the same pattern of localized domain walls but with the local σ^z magnetizations all reversed. By symmetry, the exact eigenstates in this ordered phase in the limit of large L are Schrödinger cat states of the form $|n, \pm\rangle \approx (|\phi\rangle_n \pm \hat{P}|\phi\rangle_n) / \sqrt{2}$, just like the ground states, but each such pair of excited eigenstates has its own magnetization pattern. As a result, the broken-symmetry state $|\phi\rangle_n$ has exponentially small energy uncertainty, and a system prepared in this state is metastable, with the magnetization pattern surviving for times exponentially large in system size. Equivalently, the spin autocorrelation function in any state $\langle \sigma_i^z(0) \sigma_i^z(t) \rangle$ only decays to zero on times exponentially large in system size. In this eigenstate-ordered phase, the system spontaneously breaks the Z_2 symmetry in each of its many-body eigenstates. It is evident from the above argument that the symmetry breaking is protected by the localization of the domain walls. The fact that the Landau-Peierls theorem forbids spontaneous symmetry breaking at these energies in one dimension is irrelevant because the Landau-Peierls theorem assumes thermal equilibrium and therefore does not apply in this many-body-localized phase.

To further understand the nature of this broken-symmetry phase, note that the square of the equal-time spin correlation function $\langle \sigma_i^z \sigma_{i+r}^z \rangle^2$, when evaluated in any eigenstate, shows long-range order, but $\langle \sigma_i^z \sigma_{i+r}^z \rangle$ changes sign every time we cross one of the localized domain walls. Thus, although these excited eigenstates break the Z_2 Ising symmetry, they break it in a spin-glass manner, with every excited eigenstate being characterized by its own particular σ^z magnetization pattern.

Figure 6 illustrates the above discussion. Note that as we increase the disorder strength at high energy densities, we first transition from the paramagnetic thermal phase (with only short-range spin correlations) to a many-body-localized paramagnetic phase. Here, the x components of the local magnetizations are frozen in nonthermal patterns in each eigenstate, but correlations of the $\{\sigma_i^z\}$ remain short range. As the disorder strength is further increased, there occurs a second transition to a phase with localized domain walls and thus broken Z_2 symmetry and long-range spin-glass order. Both of these aforementioned transitions are transitions in the properties of the eigenstates and thus of the system's dynamics but are not thermodynamic transitions. This phase diagram has been demonstrated using exact diagonalization data for a one-dimensional model in Reference 42. The symmetry-breaking transition within the MBL phase has also been studied using the strong-randomness renormalization group (22, 76) in References 21 and 22.

Note also the spectral transition, which occurs within the broken-symmetry MBL phase and is indicated by the dashed lines in **Figure 6**. As one crosses this spectral transition, the level statistics of the many-body spectrum changes from a Poisson distribution to a paired Poisson distribution of symmetry-related doublets composed of pairs of Schrödinger cat states. For further discussion of this spectral transition, see Reference 11.

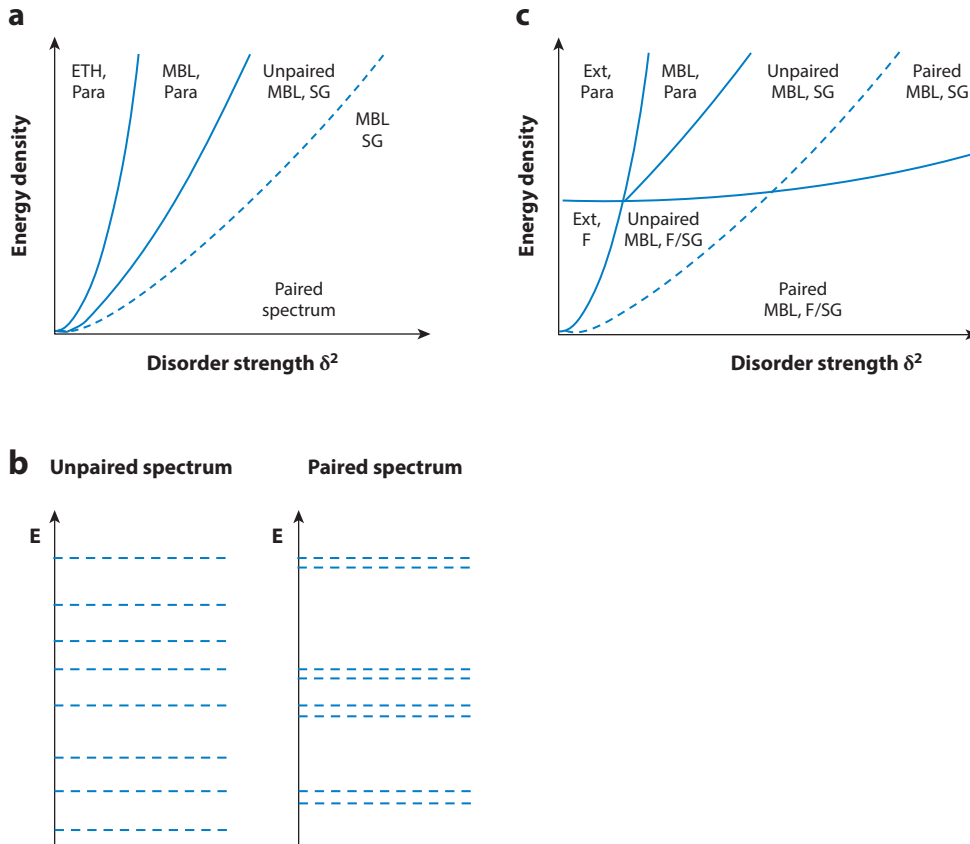


Figure 6

Schematic eigenstate phase diagrams of random quantum Ising models with ferromagnetic ground states. (a) One-dimensional system where there are no thermodynamic phase transitions. The thermal phase is labeled eigenstate thermalization hypothesis (ETH), and the localized phases are labeled many-body localization (MBL). Localization stabilizes eigenstate spin-glass (SG) order at energy densities higher than the ground state, even though this is forbidden in thermal equilibrium. (b) Spectral transition. The dashed lines in panels *a* and *c* indicate the spectral transitions between the unpaired phase, in which the energy splittings between symmetry-related Schrödinger cat eigenstates are more than the typical many-body level spacing, and the paired phase, where they are less than the typical many-body level spacing. (c) The eigenstate phase diagram in two dimensions. The principal difference from one dimension (found in panel *a*) is that ferromagnetic order (F) persists for energy densities higher than the ground state. See Reference 11 for a more detailed discussion of this figure. Figure adapted from Reference 11.

5.3. Localization-Protected Topological Order

We now discuss how localization can also protect topological order at finite energy densities, even in regimes where such order is forbidden in thermal equilibrium. We begin by illustrating this in one dimension (Equation 7). In one dimension, a Jordan-Wigner transformation turns the Ising chain (Equation 8) into a fermionic model,

$$H = -\sum_{j=1}^{L-1} iJ_j b_j a_{j+1} - \sum_{j=1}^L i h_j a_j b_j + \sum_{j=1}^{L-1} \lambda a_j b_j a_{j+1} b_{j+1}, \quad 9.$$

where the a_i and b_i are Majorana operators. For $\lambda = 0$ and in the absence of disorder $h_i = h, J_i = J$, this is simply Kitaev's Majorana chain. It is known that for $\lambda = 0$ without disorder in the regime $J > h$ there is a fermion mode that is bilocalized as two Majorana modes at the two ends of the chain and which has energy $\sim \exp(-L \ln J/h)$ (77). Thus, all eigenstates of the clean Hamiltonian, for $\lambda = 0$ and $J > h$, come in pairs, which differ by only the occupation of this fermion edge mode. Colloquially, all the eigenstates have Majorana edge modes.

In the ground state, these Majorana edge modes survive nonzero ($\lambda \neq 0$) interactions because they are protected by the bulk gap. However, in the excited states, there is no bulk gap. Indeed, in the absence of static randomness, the excited states consist of delocalized bulk excitations, which for $\lambda \neq 0$ interact with the edge Majoranas, and thus mediate interactions between the two Majorana modes at the two ends of the system. As a result, the edge Majorana modes are not robust to nonzero λ in the excited states.

In contrast, if the Dirac quasiparticles in the bulk are localized on disorder, then they cannot mediate an interaction between the two ends of the wire. Thus, we expect that in the many-body-localized phase, the excited states of the Hamiltonian (Equation 9) can also carry edge Majoranas for $\lambda \neq 0$ and thus exhibit topological order at nonzero energy density above the ground state. This is simply the dual of the statement that the Ising spin chain has spin-glass-ordered Schrödinger cat eigenstates.

A richer manifestation of topological order is found in two dimensions (11, 19). In two dimensions, the quantum Ising model has a dual description in terms of a Z_2 lattice gauge theory (78). The paramagnetic phase of the Ising model is dual to a deconfined phase of the Z_2 lattice gauge theory, with spin-flip excitations being dual to vortices (topological charges). The localization of spin flips (topological charges) on quenched disorder allows a spin-glass version of topological order to persist even in highly excited states of the lattice gauge theory. This topological order can be diagnosed by now familiar measures, such as Wilson loops, the Fredenhagen-Marcu order parameter, eigenstate degeneracy on a torus, and topological entanglement entropy (11).

In addition to symmetry breaking and topological order, in certain other systems there can exist symmetry-protected topological orders, which have recently received much attention in the theory literature (Reference 79 and references cited therein). Symmetry-protected topological order in ground states of such nonrandom systems is protected by a bulk gap (as well as by certain symmetries). It has been pointed out that just as localization can protect topological order in highly excited states, localization can also protect symmetry-protected topological order. For a fuller discussion of these issues, see References 24 and 25.

Fermionic systems with topological bands (such as the integer quantum Hall system) present an additional complication because of the well-known topological obstruction to constructing fully localized Wannier orbitals. In such systems, in the noninteracting limit, the single-particle localization length ξ_{SP} diverges near one or more critical single-particle energies $\{E_c\}$ with a critical exponent ν , as $\xi_{SP} \sim |E - E_c|^{-\nu}$. The stability of localization in such systems was recently investigated in Reference 80. It was established that arbitrarily weak short-range interactions trigger delocalization in partially filled bands at nonzero energy density if $\nu > 1/d$. Because general arguments (81, 82) constrain $\nu \geq 2/d$, it appears that MBL cannot occur in such systems. For a further discussion of these issues, see Reference 80.

6. SOME OPEN QUESTIONS

One could argue that the most important open question in MBL is: What are the experimental systems and techniques with which this phenomenon will be investigated in the laboratory? Some

ideas about this are discussed in, e.g., References 13–15, 40, 68, and 83–85. But this is a review of theoretical issues, so we now instead consider some of the interesting open theoretical questions.

6.1. Localization in Translationally Invariant Systems

One interesting class of open questions revolves around the issue of whether MBL or something like it can also occur in systems with translationally invariant Hamiltonians. Can randomness in the initial state of the system be enough to localize itself, even though the Hamiltonian is translationally invariant? If so, could quantum localization thus be relevant for, say, structural glasses, where certain degrees of freedom within a glassy material are in some sense self-localized and do not thermalize?

There are a variety of related types of models that have been considered here. One comes from considering the motion of ^3He impurities in solid ^4He (86, 87). If one assumes power-law long-range interactions between the impurities, then the hop of a single impurity atom to a neighboring lattice site changes the system's energy by some amount that depends on the positions of all other nearby impurities. Random initial locations of all the impurities are thus argued (86, 87) to localize all the impurity atoms if the hopping is weak enough. Another variation on this idea uses a generalized Bose-Hubbard model (88), with the randomness in the initial state being instead the large and random occupation numbers on each site. In this case, the authors proved an upper bound on the transport, demonstrating that the transport is, at most, nonperturbative in the hopping; they call this asymptotic localization (88).

Another class of models has two species of particles, one light and one heavy (89, 90). Given a random initial state of the positions of the heavy particles, one can then consider a Born-Oppenheimer approximation to the states of the light particles, and if the heavy particles produce enough disorder or enough constraints the light particles are localized by the (now stationary) heavy particles. For a lattice model, one can then allow hopping of the heavy particles, and argue that the transport is again, at most, nonperturbative in this hopping (90).

An open question here is whether such translationally invariant systems can have true MBL, with strictly zero DC transport at some nonzero hopping, or are they always only asymptotically localized, with some nonperturbative effect producing nonzero transport at any nonzero hopping? One possible source of nonperturbative transport is the rare regions in which the state is locally much less random and thus has mobile excitations (91; R. Nandkishore, D.A. Huse, S.L. Sondhi, unpublished results). If these rare regions themselves are mobile, they could conceivably thermalize the system. But even if such systems can only have asymptotic localization, they will still have a large regime of weak hopping in which systems of finite size in many ways appear to be localized.

Very recently, a new connection between MBL and structural glass has been drawn in Reference 92. This work studies a particular stochastic classical dynamics that is known to have a glass transition and maps the classical master equation governing these dynamics to a quantum Hamiltonian at a Rokhsar-Kivelson (RK) point (Reference 93 and references contained therein). A numerical investigation reveals that the many-body eigenstates of the quantum Hamiltonian violate the ETH both at the RK point and, apparently, for small perturbations away from it. The connection to other works on MBL in translationally invariant systems (88–91) is not presently clear.

6.2. The Many-Body Localization Phase Transition

Another set of open questions is about the nature (the universality class) of the MBL phase transition between the thermal and localized phases as the randomness is increased. The first

question is whether there is only one phase transition, or could there possibly be some sort of intermediate phase that is neither fully localized nor fully thermal (94)? Studies of this transition so far have mostly been numerical studies based on exact diagonalization of relatively short one-dimensional systems (16, 17, 42). Very little is known about this transition, so how to do a proper finite-size scaling analysis of these numerical data remains unclear. One of these studies did investigate dynamic scaling and probability distributions near the transition, suggesting dynamic critical exponent $z \rightarrow \infty$ and the possibility that the transition is governed by an infinite-randomness fixed point (17). This suggests that the transition might be treated with some version of a strong-randomness renormalization group.

Recently, the strong subadditivity of entanglement entropy has been used to establish the fact that if there is a direct continuous phase transition from a thermal phase obeying the ETH to an MBL phase, then the ETH remains true at the phase transition (94). This result constrains the possible theories of the MBL transition.

A different approach was taken in Reference 75, which examined the local spectral functions in the thermal phase in the vicinity of the localization transition. The assumption in Reference 75 was that since the local spectra are discrete in the localized phase and continuous in the thermal phase, the local spectra in a thermal phase that is close to localization should become exceedingly inhomogenous and may be characterized by a spectral line width that goes to zero as the localized phase is approached. The spectral line width, which is proportional to observables such as the DC conductivity, may then be viewed as an order parameter for the transition. The scaling of the line width with control parameters, such as disorder strength, interaction strength, and temperature, was worked out in a self-consistent mean-field approximation. However, any such mean-field theory is expected to fail in the immediate vicinity of the critical point, so a description of the critical physics remains an open problem. Progress on understanding more about this phase transition remains an outstanding challenge.

7. CONCLUSION

In this review, we have surveyed some of the present theoretical understanding of quantum thermalization and many-body Anderson localization. We emphasize that although significant progress has been made in understanding these phenomena in recent years, there still remain many open issues, and much of the present understanding is in the form of conjectures or hypotheses that need to be either verified or replaced with more accurate ideas. Thus, this is a field that is still in its infancy. This particularly applies to possible experimental investigations of MBL, where the field is still so young that it seemed to us too early to attempt a review of the possibilities. We look forward to new researchers joining in to work on these topics and thus advancing the future progress of this field.

DISCLOSURE STATEMENT

The authors are not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

ACKNOWLEDGMENTS

We thank Ehud Altman, Boris Altshuler, Ravin Bhatt, Jürg Fröhlich, Sarang Gopalakishnan, John Imbrie, Sonika Johri, Vedika Khemani, Hyungwon Kim, Joel Lebowitz, Marcus Müller, Vadim Oganesyan, Arijeet Pal, Andrew Potter, Gil Refael, Shivaji Sondhi, Tom Spencer, and Ronen Vosk for collaborations and/or discussions.

LITERATURE CITED

1. Sakurai JJ. 1985. *Modern Quantum Mechanics*. Menlo Park, CA: Benjamin/Cummings
2. Trabesinger A. 2012. *Nat. Phys. Insight* 8:263
3. Cirac JI, Zoller P. 2012. *Nat. Phys. Insight* 8:264–66
4. Bloch I, Dalibard J, Nascimbene S. 2012. *Nat. Phys. Insight* 8:267–76
5. Blatt R, Roos CF. 2012. *Nat. Phys. Insight* 8:277–84
6. Aspuru-Guzik A, Walther P. 2012. *Nat. Phys. Insight* 8:285–91
7. Houck AA, Tureci HE, Koch J. 2012. *Nat. Phys. Insight* 8:292–99
8. Kardar M. 2007. *Statistical Physics of Particles*. Cambridge, UK: Cambridge Univ. Press. 1st ed.
9. Polkovnikov A, Sengupta K, Silva A, Vengalatorre M. 2011. *Rev. Mod. Phys.* 83:863–83
10. Anderson PW. 1958. *Phys. Rev.* 109:1492–505
11. Huse DA, Nandkishore R, Oganesyan V, Pal A, Sondhi SL. 2013. *Phys. Rev. B* 88:014206
12. Fleishman L, Anderson PW. 1980. *Phys. Rev. B* 21:2366–77
13. Altshuler BL, Gefen Y, Kamenev A, Levitov LS. 1997. *Phys. Rev. Lett.* 78:2803–6
14. Gornyi IV, Mirlin AD, Polyakov DG. 2005. *Phys. Rev. Lett.* 95:206603
15. Basko DM, Aleiner IL, Altshuler BL. 2006. *Ann. Phys.* 321:1126–205
16. Oganesyan V, Huse DA. 2007. *Phys. Rev. B* 75:155111
17. Pal A, Huse DA. 2010. *Phys. Rev. B* 82:174411
18. Imbrie JZ. arXiv:1403.7837
19. Bauer B, Nayak C. 2013. *J. Stat. Mech.* 2013:P09005
20. Zangara PR, Dente AD, Lucci A, Levstein PR, Pastawski HM. 2013. *Phys. Rev. B* 88:195106
21. Vosk R, Altman E. 2013. *Phys. Rev. Lett.* 110:067204
22. Pekker D, Refael G, Altman E, Demler E, Oganesyan V. 2014. *Phys. Rev. X* 4:011052
23. Vosk R, Altman E. 2014. *Phys. Rev. Lett.* 112: 217204
24. Bahri Y, Vosk R, Altman E, Vishwanath A. arXiv:1307.4092
25. Chandran A, Khemani V, Laumann CR, Sondhi SL. 2014. *Phys. Rev. B* 89:144201
26. Serbyn M, Papic Z, Abanin DA. 2013. *Phys. Rev. Lett.* 110:260601
27. Huse DA, Oganesyan V. arXiv:1305.4915
28. Serbyn M, Papic Z, Abanin DA. 2013. *Phys. Rev. Lett.* 111:127201
29. Huse DA, Nandkishore R, Oganesyan V. arXiv:1408.4297
30. Znidaric M, Prosen T, Prelovsek P. 2008. *Phys. Rev. B* 77:064426
31. Monthus C, Garel T. 2010. *Phys. Rev. B* 81:134202
32. Berkelbach TC, Reichman DR. 2010. *Phys. Rev. B* 81:224429
33. Canovi E, Rossini D, Fazio R, Santoro GE, Silva A. 2011. *Phys. Rev. B* 83:094431
34. Canovi E, Rossini D, Fazio R, Santoro GE, Silva A. 2012. *New J. Phys.* 14:095020
35. Feigel'man MV, Ioffe LB, Mezard M. 2010. *Phys. Rev. B* 82:184534
36. De Luca A, Scardicchio A. 2013. *Europhys. Lett.* 101:37003
37. Swingle B. arXiv:1307.0507
38. Sims R, Stolz G. arXiv:1312.0577
39. Bar Lev Y, Reichman DR. 2014. *Phys. Rev. B* 89:220201
40. Aleiner IL, Altshuler BL, Shlyapnikov GV. 2010. *Nat. Phys.* 6:900–4
41. Michal VP, Altshuler BL, Shlyapnikov GV. arXiv:1402.4796
42. Kjäll JA, Bardarson JH, Pollman F. 2014. *Phys. Rev. Lett.* 113:107204
43. Khatami E, Rigol M, Relano A, Garcia-Garcia AM. 2012. *Phys. Rev. E* 85:050102
44. D'Alessio L, Polkovnikov A. 2013. *Ann. Phys.* 333:19–33
45. Ponte P, Chandran A, Papic Z, Abanin DA. arXiv:1403.6480
46. Rigol M, Dunjko V, Yurovsky V, Olshanii M. 2007. *Phys. Rev. Lett.* 98:050405
47. Cassidy AC, Clark CW, Rigol M. 2011. *Phys. Rev. Lett.* 106:140405
48. Caux JS, Essler FHL. 2013. *Phys. Rev. Lett.* 110:257203
49. Lychkovskiy O. 2013. *Phys. Rev. A* 87:022112
50. Deutsch JM. 1991. *Phys. Rev. A* 43:2046–49

51. Srednicki M. 1994. *Phys. Rev. E* 50:888–901
52. Tasaki H. 1998. *Phys. Rev. Lett.* 80:1373–76
53. Rigol M, Dunjko V, Olshanii M. 2008. *Nature* 452:854–58
54. Rigol M, Srednicki M. 2012. *Phys. Rev. Lett.* 108:110601
55. Rigol M. 2009. *Phys. Rev. Lett.* 103:100403
56. Rigol M, Santos LF. 2010. *Phys. Rev. A* 82:011604
57. Ikeda TN, Watanabe Y, Ueda M. 2011. *Phys. Rev. E* 84:021130
58. Ikeda TN, Watanabe Y, Ueda M. arXiv:1202.1965
59. Dubey S, Silvestri L, Finn J, Vinjanampathy S, Jacobs K. 2012. *Phys. Rev. E* 85:011141
60. Steinigeweg R, Herbych J, Prelovsek P. 2013. *Phys. Rev. E* 87:012118
61. De Luca A. arXiv:1302.0992
62. Beugeling W, Moessner R, Haque M. 2014. *Phys. Rev. E* 89:042112
63. Steinigeweg R, Niemeyer H, Gogolin C, Gemmer J. 2014. *Phys. Rev. Lett.* 112:130403
64. Khlebnikov S, Kruczenski M. arXiv:1312.4612
65. Mueller MP, Adlam E, Masanes L, Wiebe N. arXiv:1312.7420
66. Evers F, Mirlin AD. 2008. *Rev. Mod. Phys.* 80:1355–417
67. Iyer S, Oganessian V, Refael G, Huse DA. 2013. *Phys. Rev. B* 87:134202
68. Serbyn M, Knap M, Gopalakrishnan S, Papic Z, Yao NY, et al. arXiv:1403.0693
69. Bardarson JH, Pollman F, Moore JE. 2012. *Phys. Rev. Lett.* 109:017202
70. Kim H, Huse DA. 2013. *Phys. Rev. Lett.* 111:127205
71. Lieb EH, Robinson D. 1972. *Commun. Math. Phys.* 28:251–57
72. Bray AJ, Moore MA. 1987. *Phys. Rev. Lett.* 58:57–60
73. Nandkishore R, Gopalakrishnan S, Huse DA. 2014. *Phys. Rev. B* 90:064203
74. Johri S, Nandkishore R, Bhatt RN. arXiv:1405.5515
75. Gopalakrishnan S, Nandkishore R. arXiv:1405.1036
76. Fisher DS. 1995. *Phys. Rev. B* 51:6411–61
77. Fendley P. 2012. *J. Stat. Mech.* 2012:P11020
78. Fradkin E, Shenker SH. 1979. *Phys. Rev. D* 19:3682–97
79. Vishwanath A, Senthil T. 2013. *Phys. Rev. X* 3:011016
80. Nandkishore R, Potter AC. arXiv:1406.0847
81. Harris AB. 1974. *J. Phys. Chem.* 7:1671–92
82. Chayes JT, Chayes L, Fisher DS, Spencer T. 1986. *Phys. Rev. Lett.* 57:2999–3002
83. Basko DM, Aleiner IL, Altshuler BL. 2007. *Phys. Rev. B* 76:052203
84. Kwasigroch MP, Cooper N. 2014. *Phys. Rev. A* 90:021605(R)
85. Yao NY, Laumann CR, Gopalakrishnan S, Knap M, Mueller M, et al. arXiv:1311.7151
86. Kagan Yu, Maksimov LA. 1984. *Zh. Eksp. Teor. Fiz.* 87:348–65
87. Kagan Yu, Maksimov LA. 1984. *Sov. Phys. JETP* 60(1):201–10
88. Huveneers F, De Roeck W. arXiv:1308.6263
89. Grover T, Fisher MPA. arXiv:1307.2288
90. Schiulaz M, Muller M. 2014. *AIP Conf. Proc.* 1610:11
91. Huveneers F, De Roeck W. arXiv:1405.3279
92. Hickey JM, Genway S, Garrahan JP. arXiv:1405.5780
93. Castelnuovo C, Chamon C, Mudry C, Pujol P. 2005. *Ann. Phys.* 318(2):306–44
94. Grover T. arXiv:1405.1471



Contents

Innovations in Statistical Physics <i>Leo P. Kadanoff</i>	1
Many-Body Localization and Thermalization in Quantum Statistical Mechanics <i>Rahul Nandkishore and David A. Huse</i>	15
Composite Fermion Theory of Exotic Fractional Quantum Hall Effect <i>Jainendra K. Jain</i>	39
The Statistical Physics of Athermal Materials <i>Dapeng Bi, Silke Henkes, Karen E. Daniels, and Bulbul Chakraborty</i>	63
Napoleon Is in Equilibrium <i>Rob Phillips</i>	85
Assembly of Biological Nanostructures: Isotropic and Liquid Crystalline Phases of Neurofilament Hydrogels <i>Cyrus R. Safinya, Joanna Deek, Roy Beck, Jayna B. Jones, and Youli Li</i>	113
Plutonium-Based Heavy-Fermion Systems <i>E.D. Bauer and J.D. Thompson</i>	137
Exciton-Polariton Bose-Einstein Condensates <i>Benoît Deveaud</i>	155
Marginal Stability in Structural, Spin, and Electron Glasses <i>Markus Müller and Matthieu Wyart</i>	177
Ultracold Atoms Out of Equilibrium <i>Tim Langen, Remi Geiger, and Jörg Schmiedmayer</i>	201
Motility-Induced Phase Separation <i>Michael E. Cates and Julien Tailleur</i>	219

Physics of Viral Shells <i>Robijn F. Bruinsma and William S. Klug</i>	245
Amplitude/Higgs Modes in Condensed Matter Physics <i>David Pekker and C.M. Varma</i>	269
Symmetry-Protected Topological Phases of Quantum Matter <i>T. Senthil</i>	299
Spatial Localization in Dissipative Systems <i>E. Knobloch</i>	325
Topological Crystalline Insulators and Topological Superconductors: From Concepts to Materials <i>Yoichi Ando and Liang Fu</i>	361
Universal Dynamics and Renormalization in Many-Body-Localized Systems <i>Ehud Altman and Ronen Vosk</i>	383
Quantum Oscillations in Hole-Doped Cuprates <i>Suchitra E. Sebastian and Cyril Proust</i>	411

Errata

An online log of corrections to *Annual Review of Condensed Matter Physics* articles may be found at <http://www.annualreviews.org/errata/conmatphys>