

Many-body Quantum Dynamics

Lecture I

10 October 2023

Abstract

We give an overview over the contents of our lecture course. Part of it is going to be provision of background, but more importantly we would like to provide some motivation for the questions that we are going to be discussing. Much of what we're going to say today you will already know one way or another. But we'd like to focus on the fact that some of the connections between different topics have not been known until quite recently. Indeed, one of the great advances the field has known is that we understand how different pictures are interlinked, and much of our lecture course will be devoted to the results which have made these insights possible. As we go along, we will be flagging the topics that will appear in various lectures of this course.

1 Introduction

So initially we address the two items: what do we know? And how do the things we know fit together?

Let us start with equilibrium physics. Here we have two conceptually largely separate pictures, namely equilibrium thermodynamics on one hand, and statistical mechanics on the other hand.

Beyond equilibrium, the landscape is much more varied and poorly understood and there is less systematic understanding to guide us. Topics here include quenches, periodically driven systems, and even fully aperiodic unitary quantum evolution.

There is also the topic of open quantum systems, which however is not so central to this course, but which will likely be covered in a lecture in January.

At the same time, the basic formulation of quantum dynamics in terms of Schrödinger's equation is not wedded either to equilibrium and on equilibrium. Therefore, there should be some simplifications when using Schrödinger's equation to describe systems in thermal equilibrium. Understanding which form these take, and what insights follow, is the sense that we will be discussing the role of quantum many body dynamics in this lecture course.

2 Different conceptual frameworks for many-body physics

2.1 Thermodynamics

So, let us start with what we have known for the longest: thermodynamics is a theory which was largely complete by the 19th century. It is entirely non-microscopic in the sense that it can be formulated without any recourse to microscopic information – a fortiori, it was developed before the microscopic nature of matter had even been elucidated.

This is a point that is extremely important to appreciate. While we nowadays often teach thermodynamics as a consequence of statistical mechanics, it is actually conceptually largely

independent, and much of its power in fact derives from the fact that it does not presuppose any microscopic knowledge of the systems it studies. In some sense, thermodynamics tells you what the minimal amount of information is that you need to describe systems in equilibrium.

Two central concepts are those of equilibration and thermalisation. Equilibration states that a system left to its own devices will eventually reach a steady state. Thermalisation, in turn, posits that the nature of this steady state depends on only very few parameters – particle number/chemical potential, energy/temperature, volume/pressure.

Note that these insights likely predate the scientific method itself – presumably, our cave dwelling ancestors will have had an appreciation of a lake being predictably frozen when the temperature has been low for long enough. Nonetheless, thermodynamics has built a solid edifice, based on the ideas of quasistatic processes, equilibrium states, as well as **thermodynamic potentials**. (As an aside, notice that the term 'adiabatic' has incompatible meanings in thermodynamics and quantum dynamics: in the latter, it describes a process which is so slow that no transitions between levels take place; while in the former, it refers to processes in which no heat is exchanged with the outside world).

Basically, what thermodynamics tells you is to maximise the entropy of a systems subject to the constraints imposed on it. The nature of the constraints determines the natural variables in the setting under consideration, from which it follows which thermodynamic potential is extremised in equilibrium. Further, derivatives of the thermodynamic potential with respect to its natural variables yields the thermodynamic behaviour of the system in question, and Maxwell relations provide links and consistency relations between these.

One set of questions that will be covered in these lectures is: what are alternatives to equilibrium thermodynamics, and what are their properties? This is a priori a large field, and much of the world around us is not in equilibrium. As physicists, we start with the simplest deviations from equilibrium, and we will be discussing *many-body localisation* as an example of a system which fails to equilibrate; *driven systems* as an instance of a setting where equilibration is precluded by the absence of a static Hamiltonian, with *quantum circuits* describing a situation where the idea of Hamiltonian evolution is largely discarded from the start.

As a sign of how powerful thermodynamics is, it is worth recalling that it made some robust but demonstrably wrong predictions which led the way to the discovery of quantum mechanics—most saliently, in the specific heat of solids, and the properties of blackbody radiation. There, predictions based on classical equipartition yielded incorrect functional dependencies of quantities such as the specific heat as a function of temperature, $C(T)$, by failing to take into consideration the correct quantum occupation numbers of the modes in question.

2.2 Statistical mechanics

The central quantity in statistical mechanics is the partition function

$$Z = \sum_i \exp\left(-\frac{E_i}{k_B T}\right) \quad (1)$$

where the sum runs over all microstates of the system, each of which contributes its Boltzmann factor, where k_B is Boltzmann's constant and $T = 1/(k_B \beta)$ is the temperature.

The thing to notice immediately is the enormous amount of information required here—for a general system, determining (indeed, even enumerating) all microstates can be a non-trivial, and often intractable, exercise.

In modern courses on thermal physics, thermodynamics is often, very ahistorically, derived from statistical mechanics by evaluating thermodynamic potentials from the partition function. That this should be possible is not too surprising, given how much more information goes into the partition function compared to thermodynamics. The minimisation of the free energy,

$\mathcal{F} = \mathcal{U} - TS$, e.g., comes about simply by considering that the number of states at a given energy is the exponential of the entropy, $\log \rho_E = S(E)$, so that one can rewrite

$$Z = \sum_i \exp\left(-\frac{E_i}{k_B T}\right) = \int_{E_{\min}}^{E_{\max}} dE \rho(E) \exp(-\beta E) = \int_{E_{\min}}^{E_{\max}} dE \exp(-\beta E + S) = \exp(-\beta \mathcal{F}) \quad (2)$$

with the last equality holding if the argument of the exponential is sufficiently strongly peaked.

Much thought has gone into the description of the process of discarding the information of statistical mechanics, which goes under names like coarse-graining (such as block-spin transformation), the derivation of effective actions, RG flows etc. Also, the description of statistical mechanics permits to evaluate ‘mesoscopic’ deviations from the behaviour of thermodynamically large systems.

2.3 Quantum mechanics

Perhaps the most fundamental formulation of quantum mechanics proceeds by writing down the Schrödinger equation

$$i\hbar \frac{d}{dt} |\psi(t)\rangle = \hat{H} |\psi(t)\rangle . \quad (3)$$

This is yet a totally different language. What appears now is a wavefunction living in Hilbert space, and the Hamiltonian, a hermitian operator; moreover, the equation describes the evolution of a system in real time, while thermodynamics and statistical mechanics had no obvious consideration of time-dependence—indeed, thermodynamics with its quasi-static processes and equilibration/thermalisation at long times would seem to be formulated with an entirely ‘opposite’ attitude.

In some ways, Schrödinger’s equation feels much more closely linked to statistical mechanics. Its eigenfunctions can be put into the expression for the partition function, and in this way, classical and quantum mechanics can be linked with a minimal amount of fuss.

But if we want to understand how these ways of describing many-body systems are interlinked, we need to sharpen our terminology somewhat.

3 What is quantum many-body dynamics

It is worth emphasizing at the very outset that the meaning of words as elementary as quantum and many-body have a strong community (and personal) dependence. Even if you listen to a presentation on the subject of quantum chaos, you will find that different subcommunities have a conflicting vocabulary that seems almost designed to create misunderstandings. The aim in the following is to discuss the issues in question, and to introduce the meaning these words have in our context, rather than to provide a dictionary between different usages.

3.1 Many-body

Many-body physics can be taken to be a complement to single-particle, few-particle, or even mesoscopic physics. It can also be used to delineate systems of non-interacting particles, as their effective description can be cast in terms of single-particle physics.

So, basically, in many-body physics, we study the behaviour of large numbers (e.g., Avogadro’s number, $N_A \sim 6 \times 10^{23}$) of interacting constituent elements.

3.1.1 Generic many-body systems

It quickly becomes tedious to list all features/bugs that a system should not exhibit to fall under the heading of many-body system. One therefore uses the shorthand 'generic' many-body system to delineate wanted from unwanted properties – and these may differ depending on context. One e.g. often excludes integrable systems (with an infinite number of local integrals of motion). Other features, such as local interactions, are often also at least implied.

At this stage, it is perhaps more helpful to give a concrete example of a many-body system. The one we choose is a chain of spins. Most simply, these spins have two states, pointing either up or down. The total configuration space for a system of N spins then has size 2^N . Obviously, this quickly becomes an astronomically large number, even compared to astronomically large numbers such as the $\approx 10^{80}$ atoms (give or take a few orders of magnitude) in our universe.

Note that a lattice system like this has both a discrete set of basis states, as well as being defined on a discrete set of sites in real space. Often, not having unbounded/continuous spaces is desirable. The physical content of this demand is not entirely clear, e.g. we do not know if space actually is continuous (on, say, the Planck scale $l_p = \sqrt{\hbar G/c^3}$).

Also, note that in quantum mechanics, continuously variable superpositions of discrete basis states can continuously parametrise expectation values. Just think of a spinor

$$\cos\left(\frac{\theta}{2}\right)\begin{pmatrix} 1 \\ 0 \end{pmatrix} + \sin\left(\frac{\theta}{2}\right)\begin{pmatrix} 0 \\ 1 \end{pmatrix} \quad (4)$$

which has the expectation value of the z-component of a $S = 1/2$ vary from the north to the south pole as θ is increased.

What these discreteness requirements achieve is to have bounded energy spectra, e.g. for the kinetic energy of a mobile particle hopping with matrix element $-t$ between adjacent sites on a chain, whose dispersion relation is $\epsilon(k) = -2t \cos k$.

3.2 Quantum

Discussions of what the word 'quantum' mean often descend into endless philosophising. This is because, even a century after its formulation, there still is no agreement on what is the essence, or rather, minimal set of defining characteristics, of quantum mechanics. Let us sidestep this issue and continue with our example of a $S = 1/2$ spin chain. At the end of this lecture, we'll say a few more words about this issue from various perspectives.

3.2.1 Interactions: spin-1/2 chain

The classical configuration space can also be thought of as a basis of Hilbert space–this is often referred to as the computational basis.

The classical energy function, e.g., can involve only the z -components, i.e. the Ising spins $S_i^z = \pm 1$:

$$E(\{S_k\}) = \sum_k (J_1 S_k^z \cdot S_{k+1}^z + J_2 S_k^z \cdot S_{k+2}^z) . \quad (5)$$

$J_{1,2}$ are exchange constants, positive being antiferromagnetic, negative ferromagnetic.

This can at the same time be thought of as a statement about operators–the Ising spins now becoming Pauli matrices

$$\sigma^z = \begin{pmatrix} +1 & 0 \\ 0 & -1 \end{pmatrix} \quad (6)$$

and the Hamiltonian is diagonal already in this basis. In this sense, the classical configurations become quantum wavefunctions, in particular eigenfunctions of the Hamiltonian. These form a complete basis set, as the Hamiltonian is hermitian.

This is a highly non-generic Hamiltonian, as one might expect: relabelling a classical system a quantum one is not generic. This can be made quantitative by looking at what is known as the level statistics: the spectrum of this Hamiltonian contains huge degeneracies—the energy depends only on the number of frustrated bonds, and their location/separation is immaterial.

In generic systems, such degeneracies are unstable, and the probability of finding two closely spaced levels is suppressed compared to what one would find if the eigenenergies were drawn at random in a given interval, according to the local density of states. This subject will be covered under the heading of *random matrix theory*.

It is now only a small step towards a generic quantum Hamiltonian. By adding an off-diagonal term to the Hamiltonian, the level statistics is restored, and the degeneracies are lifted. This is achieved by

$$H_{\text{TFIM}} = \sum_k (J_1 S_k^z \cdot S_{k+1}^z + J_2 S_k^z \cdot S_{k+2}^z) + \Gamma \sum_k S_k^x \quad (7)$$

with the x-component now given by the appropriate Pauli matrix

$$\sigma^x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \quad (8)$$

This is known as the transverse field Ising model, one of the paradigmatic models of many-body physics, the workhorse of many studies from quantum phase transitions to out-of-time-ordered correlators. The nearest-neighbour version, $J_2 = 0$, was first studied by Roger Elliott in Oxford about 50 years ago, and this is exactly soluble and hence non-generic. With nonzero J_2 , it is fully generic (although the crossover between the two behaviours is quite interesting; this will be covered under the heading of *prethermalisation*).

For the studies of this kind of spin Hamiltonian throughout this course, we'll be using the QuSpin package developed by Marin Bukov and collaborators.

3.2.2 What is an eigenstate (good for)?

Let us now return to the question of levels and their separation. The total bandwidth, i.e. difference between the maximal and minimal energy of the system, is proportional to the system size, $N = L$ in one dimension. However, the number of levels grows exponentially, 2^N . The mean spacing between levels is therefore exponentially small in N :

$$\Delta_E \sim N 2^{-N}. \quad (9)$$

(This is in fact not true uniformly throughout the spectrum. Near its edges, i.e. the ground state and most excited state, the levels are much more sparse. This can go along with an 'emergent integrability' of quasiparticles. This kind of physics – nature of ground states, and excitations around them – is not an important part of this lecture course, even though it is obviously one of the central quests of condensed matter physics. At any rate, several lecture courses here in Dresden and elsewhere are devoted to this subject).

In this sense, we are going to be studying finite-temperature properties, which means that there is a finite energy density. If there is a quasiparticle description, this amounts to a finite density of quasiparticles. Given the abovementioned emergent integrability, the generic behaviour may only emerge on very large length- and timescales.

How relevant is such an eigenstate for practical purposes? There are two items to note here.

First, there is a timescale related to Δ_E , as given by one of Heisenberg's uncertainty relations:

$$\Delta_E \delta_t \sim \hbar. \quad (10)$$

What this says is that, in order to construct a state with an energy defined sharply enough that it is separated from the neighbouring eigenstate, one needs a time which scales exponentially in system size: $\delta_t \sim 2^N$. Again, this timescale is astronomical, and will in particular not be accessible to coherent quantum experiments. In this sense, one can question the very existence of generic eigenstates.

The second item is the description of the eigenstate. It lives in Hilbert space, the vector space spanned by the 2^N computational basis states. To specify a general vector in this space thus requires exponentially many complex coefficients. It is easy to convince yourself that it is impossible to write down this amount of information on any classical storage medium, even using all the resources available in the universe.

Thus, we can not only not prepare such a state, we can also not write it down. (There are special states for which this is not true. Most spectacularly, the ground states of local Hamiltonians are considerably less complex than generic vectors in Hilbert space. The DMRG revolution of the last two decades has been driven by the possibility of writing these states efficiently using so-called matrix product states.)

So, where does the information go, and why might it still be useful to think of properties of generic eigenstates? This is going to be the subject of the lectures on *eigenstate thermalisation*.

3.3 Actual quantum dynamics

So, what about actual quantum dynamics? Essentially everything discussed so far is related to statics, and one may wonder what happens to actual two-time correlation functions $\langle S(t)S(0) \rangle$? Very similar issues as in the case of the eigenstates arise here. To see this, consider the Lehmann representation

$$\langle S(t)S(0) \rangle = \frac{1}{Z} \text{Tr}_{\{S\}} \exp(-\beta H) S_a(t) S_b(0) = \frac{1}{Z} \sum_i \exp(-\beta E_i) \langle i | S(t) S(0) | i \rangle \quad (11)$$

$$= \frac{1}{Z} \sum_{i,j} \exp(-\beta E_i) \langle i | \exp(iHt) S_a(0) \exp(-iHt) | j \rangle \langle j | S_b(0) | i \rangle \quad (12)$$

$$= \frac{1}{Z} \sum_{i,j} \exp(-\beta E_i) \exp(it(E_i - E_j)) \langle i | S_a(0) | j \rangle \langle j | S_b(0) | i \rangle . \quad (13)$$

We therefore need not only the eigenstates, but all the matrix elements between them of the operators in question. This, again, is going to be subject of the eigenstate thermalisation.

This raises a very natural question: where does all the information encoded in the precise eigenfunctions go, as a matter of principle? This subject is covered under the heading of *information scrambling*.

3.3.1 Heisenberg vs. Schrödinger picture, and operator growth

In quantum mechanics courses, it is often stated that there are two alternative pictures of quantum mechanics. In one, the Schrödinger picture, it is the wavefunctions which are time-dependent, and the operators are stationary. In the other, Heisenberg's, it is the operators which time-evolve, while the wavefunctions remain stationary, with a time-evolution (in the absence of an explicit time-dependence of the operators) given by

$$-i\hbar \frac{dS}{dt} = [H, S] \quad (14)$$

In the above Lehmann representation, the Heisenberg picture was used, while initially, we wrote down the Schrödinger equation.

In some sense, the Heisenberg picture is more encompassing, as the Schrödinger equation a priori deals with the evolution of a single wavefunction, and not the full operator: the dimension of the latter is roughly the square of the dimension of the former. Of course, ultimately the two descriptions are fully equivalent.

Considerations of how operators evolve has yielded plenty of insights into the structure of the above subjects. This will be discussed in a lecture under the heading of *operator spreading*, which is closely related to information scrambling mentioned above.

4 Beyond equilibrium

Non-equilibrium physics is a negative definition – what it's not – a little bit like non-Fermi liquids, or non-bananas: most of the world is in that respective state. The question then in any given setting is: what goes out of the window compared to the equilibrium case?

So one of the things that we are going to be interested in is to what extent we can still use the framework of equilibrium thermodynamics for describing systems beyond their equilibrium properties. One particular element we can focus on here is the notion of order and disorder, and the transitions separating different phases of matter. More precisely can there even be different phases outside of equilibrium? And, can there possibly be phases beyond equilibrium, which do not exist at all in equilibrium. This will lead us to the study of *time crystals*, which represent a spatiotemporal form of order.

A mild way of deviating from stationary Hamiltonians is to add a periodic time-dependence, so that only a single frequency and its harmonics appear in the time-dependence. This is the realm of *Floquet physics*, and of *Floquet engineering*, which plays an important role in experiments in cold atoms but above all has turned out to be hugely productive in terms of advancing our conceptual understanding of non-equilibrium physics.

5 Considerations on the classical-quantum dichotomy

Let us conclude by saying a few words about what we actually mean by quantum mechanics. This is a question that may at first seem either very simple or very philosophical. However answering it requires us to address if you fundamental issues about how we think about physics.

A few subquestions here are: In what sense are we describing a system when we write down Schrödinger's equation? And, what do we actually mean by a system being quantum mechanical in specific distinction to being classical? Which phenomena (such as entanglement or interference) are intrinsically quantum mechanical?

The latter question focuses on why it is so hard to draw a clear distinction. Wave phenomena and interference are commonly encountered in classical mechanics, and whether or not a state is a result of interference processes is to some extent basis-dependent. Also, various effectively classical descriptions of quantum mechanical systems exist. One of my colleagues once opined that anything we understand is, almost by definition, ultimately classical.

So, let us perhaps start from a different direction. Schrödinger's equation is in effect a low energy theory. For the time being, we are lacking a complete quantum mechanical description of the universe, i.e. a unification of all the forces etc. While this is the case, Schrödinger's equation has the status of a low energy theory; quantum mechanical descriptions at higher energies may involve relativistic quantum mechanics etc, which are to some extent taken into account in e.g. spin-orbit coupled models.

In practise, while the general form of Schrödinger's equation quoted above is always the same, it can take on a wide variety of guises, even mimicking relativistic theories, such as in

the description of electrons in graphene.

How does classical mechanics fit in here? Ultimately, everything is quantum mechanical, and classical mechanics is an emergent property of some quantum systems. Formally, this often occurs via a singular limit: scrambling gets worse for larger spins with their larger Hilbert spaces, but physically, larger spins are better approximated semi-classically. How this occurs remains in part an open question.

As described above, if you have a simple classical Ising model, you can turn it into a quantum using models by simply adding a term for the Hamiltonian such as a transverse field whose action in the computational basis is to flip a spin. In this sense, the difference between a classical and a quantum model rest solely in the way its dynamics is described, and the classical setting which has no dynamics at all is in this sentence a limiting case of the quantum dynamics.

One very topical and concrete example is the case of a quantum computer. The challenge is, from a commercial perspective, to demonstrate that a quantum computers can do things that classical computers cannot. This is known as quantum supremacy or quantum advantage. Companies such as Google are betting a lot of money on being able to pull off the achievement of having a quantum platform which is superior to a classical one.

Therefore one tries to identify a situation where an operation takes place which cannot be achieved simply by having a classical system do the computation. To understand what is meant by a computation, let us turn to the Church-Turing hypothesis: “Every computation that can be carried out in the real world can be effectively performed by a Turing Machine.” So, a quantum computer would be one which cannot be efficiently emulated by a classical computer.

The way quantum complexity poses this question is in the form of a black box experiment. Can I convince you that a given relation between input an output cannot be obtained by purely classical means. This is not an easy question to answer.

An enormous cottage industry is devoted to this aspect of quantum information theory, devising various measures of entanglement and non-classical correlations. This is related to Alain Aspect’s Nobel Prize for experiments studying the nature of quantum entanglement.

In the church-Turing sense, a quantum computer would represent the most general quantum time evolution, exemplified by the quantum circuit that execute. And it is important to emphasize how far quantum computers are from reaching us at the moment, and a present, it seems more appropriate to think of quantum computers as very coherent many body systems, rather than universal computational machines.

It is in the sense that the study of generic *quantum circuits* has taken up an additional level of interest, and these will be covered in a lecture of their own.