10
Topology out of equilibrium

When we say we consider the ground states of closed quantum systems, or their equilibrium finite-temperature properties, this implies considerable restrictions on the kind of behaviour we can expect to encounter. For instance, the ground state at zero temperature, being an eigenstate, is stationary under the evolution generated by the system Hamiltonian, itself taken to be time-independent; while at nonzero temperature, considerations of entropy maximisation proscribe certain types of ordering, as exemplified by the impossibility, encoded in the Mermin-Wagner theorem, of breaking continuous symmetries in $d = 2$.

While there is, more or less, one type of thermal equilibirum, many non-equilibrium settings are possible. For example, by considering an initial state which is not an eigenstate of the system Hamiltonian, one obtains a time dependent problem even for a time-independent Hamiltonian; this has become to be known as a quantum quench. If the Hamiltonian itself is time dependent, one refers to a driven system.

For the finite-temperature case, the implicit assumption is that a concept of temperature is defined in the first place, and that the system thermalises, i.e., reaches an effectively time-independent (as far as local observables are concerned) steady state, the properties of which depend only on a small number of parameters, such as the energy density of the initial state. However, there are cases where these assumptions are not satisfied; these include glasses and localised, see Chapter 8, systems. In particular the many-body localisation, briefly introduced below, furnishes a generic route to non-thermalisation.

In the following, we cover material which touches on each of the above items. This chapter will necessitate a fair amount of background material, which we provide as we go along. This includes sections on thermalisation and the lack thereof; the description of periodically driven (Floquet) systems, as well as Floquet engineering; and the possibility of defining phase structure out of equilibrium via the notion of eigenstate order.

Nonetheless, the non-equilibrium behaviour of matter is such a broad, complex and rich field – as the reader can verify by taking a quick look out of the window – that it is not yet possible to write anything approaching as comprehensive an account as can be done in the equilibrium setting: at present, we quite simply even

lack a similarly systematic framework. This statement holds true both for topological aspects and more broadly. However, a focus on periodic driving creates a useful structure for non-equilibrium systems loosely analogous to how assuming perfect crystals simplifies equilibrium topology via Bloch's theorem. However, despite the formal similarity between Floquet and Bloch structures, there are numerous differences in how they apply in realistic physical systems, as we start to see in the following section.

10.1 Time-dependent, and time-periodic (Floquet) Hamiltonians

The aim of this section is to follow an avenue in which a relatively gentle deviation from the equilibrium setting has proven still to be tractable while yielding qualitatively new phenomena. We first provide a compendium of simple facts about unitary time evolution generated by time dependent Hamiltonians. This will serve as a springboard for the special case of when this time-dependence is periodic. Such systems are known as Floquet systems, and we will find that they host a number of interesting new topological phenomena which are beyond the reach of static Hamiltonians.

The formal solution of the time-dependent Schrödinger equation

$$i\hbar \frac{d|\psi\rangle}{dt} = H(t)|\psi\rangle \tag{10.1}$$

is given by the unitary time evolution operator

$$\mathcal{U}(t, t_0) = \mathcal{T} \exp\left[-\frac{i}{\hbar} \int_{t_0}^{t} dt\, H(t)\right] . \tag{10.2}$$

This can in turn be used to implicitly define an *effective Hamiltonian*,

$$\exp\left[-\frac{i}{\hbar}(t - t_0)H_{\text{eff}}\right] := \mathcal{U}(t, t_0) . \tag{10.3}$$

At this stage, this is simply a formal definition, and we will in particular have to return to the issue of the non–single-valuedness of the logarithm. However, the value of this is that it indicates the possibility of transferring a lot of the intuition we have from regular Hamiltonians. In particular, it is guaranteed that there is a complete set of orthonormal states for the Hilbert space of the system, ensuring that any state $|\psi(t)\rangle$ can be written as a linear combination of eigenstates of the time-evolution operator.

However, at this stage, it is then natural to pose the question what this setting can possibly achieve that cannot also be achieved by, say, a time-independent Hamiltonan H_{eff}.

The answer is that a Hamiltonian thus defined does not inherit all physically important properties from the instantaneous Hamiltonians $H(t)$ on which it depends. Perhaps most importantly, H_{eff} need not be, and will in general not be, local, even

if $H(t)$ is local for all times t. (Loosely spekaing, local means involving only combinations of operators which are nearby in real space, like exchange interactions between spins at most a few lattice spacings apart.) Since several properties of many-body systems which we take for granted actually depend on the locality of Hamiltonians, systems described by a non-local H_{eff} can thus exhibit unexpected (and novel) behaviour. The time crystal described in detail below is a case in point.

10.2 Floquet basics

In the case of periodically driven systems, $H(t + T) \equiv H(t)$, the effective Hamiltonian is generally referred to as the Floquet Hamiltonian, H_F, and $\mathcal{U}_F = \mathcal{U}(0, T) = \exp\left[-\frac{i}{\hbar} T H_{\text{F}}\right]$ is the Floquet unitary.[1]

10.2.1 Discrete time translation symmetry

Floquet systems have considerable additional structure compared to the general time-dependent case. While the latter have discarded the invariance of the static case with respect to infinitesimal time translations, the Floquet problem retains a symmetry with respect to *discrete* time translations. This leads to energy conservation, which follows from Noether's theorem, to be replaced by quasi-energy conservation: quasi-energies are only defined modulo $\hbar\Omega = 2\pi\hbar/T$.

This is completely analogous to the vestigial conservation of crystal momentum in a periodic potential, as opposed to momentum conservation in the continuum. There, physically distinct momenta of a Bravais lattice are restricted to a Brioullin zone. For the case of a chain of lattice constant a, the allowed crystal momenta thus lie in the interval $[-\pi/a, \pi/a)$, with momenta differing by a reciprocal lattice vector $2\pi/a$ being equivalent. For a periodically driven system, quasienergies are similarly restricted to lie in a Floquet Brioullin zone ranging from $-\hbar\pi/T$ to $\hbar\pi/T$. The multivaluedness of the logarithm in the above definition of the effective Hamiltonian is related to the choice of Brioullin zone, with choices differing by addition of a 'reciprocal lattice vector' $2\hbar\pi/T$.

10.2.2 Floquet ensembles

In a non-equilibrium setting, the familiar constraints imposed by equilibrium thermodynamics need to be rethought: concepts like thermodynamic potentials or temperature need no longer be useful or even exist. In Box. 10.1, we give a very brief summary of the issues involved in order to make this treatment self-contained, but which fails to do justice to the rich and interesting field of non-equilibrium quantum dynamics.

---

[1] One is free to fix the 'gauge choice' of the time of the 'beginning' of the period, t_0, a point which will not be important in what follows.

Let us first address the issue of temperature, whose existence is related to energy conservation, in the same way that a chemical potential is defined only in the presence of particle number conservation. By its very nature, this item has been abandoned in the Floquet setting, as quanta of energy $2\pi\hbar/T$ can be added or subtracted, and hence there is no concept of temperature. Staying within the framework of thermodynamics, what we now need to do is maximise entropy without this constraint – which effectively means giving each state the same weight. This is also known colloquially as an 'infinite temperature ensemble', as the concomitant Boltzmann factors $\exp[-E/(k_B\,T)]$ also become state-independent when one sets $T = \infty$.

This turns out to be the generic setting for Floquet systems. It is known simply as Floquet-ETH, in analogy to the static systems obeying eigenstate thermalisation. In the Floquet case, there is obviously nothing to talk about in terms of non-trivial correlations in the long-time limit.

However, energy conservation need not be the only constraint present in a system. There can also be particle number conservation, global symmetries like a U(1) spin symmetry or, in an integrable system, any number of other constants of motion. The most straightforward way of obtaining an integrable system in the Floquet setting is to consider a set of free fermions, each of which is subject to a periodic drive, but which remain non-interacting. This setting allows importing ideas from the corresponding static setting, in particular the notion of a generalised Gibbs ensemble, see Section 10.2.3. This programme can be carried out entirely analogously for the Floquet case, where it has been christened the (Floquet-) periodic Gibbs ensemble, Floquet-PGE.

A more involved way of avoiding equilibration to an infinite temperature Floquet-ETH ensemble involves avoiding the process of equilibration altogether. This can be achieved by adding disorder to the system, so that the system ceases to be ergodic. It turns out that this can be arranged for in a way which is generic, that is to say which – unlike the integrable case mentioned above – is stable to any small perturbation of the Hamiltonian. The underlying phenomenon is known as many-body localisation (Nandkishore and Huse, 2015). We will discuss this option after covering the integrable cases in the next sections.

Box 10.1: Phase structure in and out of equilibrium

The aim of this section is to explain how one can generalise the notions of phases, and transitions between them, beyond the familiar setting of equilibrium thermodynamics. To do this, we first need to illuminate the connection between quantum many-body physics and equilibrium thermodynamics, where the notion of eigenstate thermalisation plays a central role. We then explain

---

Box 10.1: Phase structure in and out of equilibrium

---

how the phenomenon of many-body localisation (MBL) presents an alternative to thermalisation. And finally, how MBL allows for the identification of non-equilibrium phases in a crisp way.

10.2.3 Equilibration and thermalisation–and absence thereof

Basic thermodynamics is built on the twin concepts of equilibration and thermalisation. The former states that system, left to its own devices, will eventually reach a time-independent steady state. The latter implies that this steady state is determined by only a small number of parameters – such as conserved quantities like energy or particle density (or temperature/chemical potential, depending on the choice of ensemble).

Eigenstate thermalisation

This is a far cry from the microscopic picture of quantum mechanics embodied by the Schrödinger equation, which is a fully microscopic theory in which a general wavefunction is determined by the amplitudes of the different basis states it contains. For a lattice system of N spins-1/2, there are exponentially many, 2^N, of these. This is unimaginably far more information than is encoded in the thermodynamic description. Also, it is in fact impossible to prepare an exact generic eigenstate at a finite energy density above the ground state, since the adjacent levels are only an energy of order $O(2^{-N})$ away, so that Heisenberg's uncertainty principle states it would take a time $O(2^N)$ to prepare them. For a macroscopic N of a thermodynamically large sample, this time is beyond our lifetimes, if not that of the universe.

The resolution is provided by the eigenstate thermalisation hypothesis (ETH), which essentially states that local observables and correlators in generic eigenstates take on the values characteristic of thermodynamic equilibrium at the energy/particle density of the eigenstate under consideration. Then, it is no longer necessary to know all the basis state amplitudes; nor indeed is it necessary to have an exact eigenstate: any combination of quantum state with the same energy/particle densities will do.

Generalised Gibbs ensemble

It is possible to increase the number of conserved quantities yet further. There may for example also be symmetries leading to the conservation of spin or momentum density. Indeed, in integrable systems, there may be an extensive

---
Box 10.1: Phase structure in and out of equilibrium
---

number of conserved quantities.[a] Each of these then enforces its own constraint. This leads to what is known as a generalised Gibbs ensemble (GGE): each conserved quantity leads to a 'Lagrange multiplier' in the way that energy/particle density lead to the notions of temperature and chemical potential. These integrable systems are fine-tuned, in the sense that generic perturbations typically destroy the supernumerary conservation laws, and collapse the GGE to a standard thermodynamic Gibbs ensemble. The counterpart of the generalised Gibbs ensemble in Floquet systems with additional conservation laws is known as Floquet periodic Gibbs sensemble (PGE).

10.2.4 Eigenstate thermalisation, phase transitions and order

The notion of phases and phase transitions then transfers neatly from thermodynamics to eigenstates. Consider a setting with a disordered high-temperature state and a low-temperature ordered one, such as in a transverse field Ising ferromagnet in dimension $d \geq 2$ (Eq. 10.4 with couplings independent of j). In that case, generic eigenstates at high energy density do not exhibit long-range ferromagnetic order, while those at low energy do. The critical energy corresponds to the energy density of the system at the temperature of the thermodynamic phase transition.

This allows to make a connection to the quantum quenches mentioned in the introduction. Consider evolving an initial state, $|\psi_0\rangle$ with such a Hamiltonian (generally not an eigenstate thereof). Whether expectation values at long times will exhibit long-range order or not then depends only on the energy density of the initial state, i.e. the expectation value $\langle\psi_0|H|\psi_0\rangle$.

10.2.5 Many-body localisation

We next present MBL as an alternative to thermalisation in a generic many-body system; that is to say, unlike Anderson localisation (discussed in Box 8.1), MBL does not require any specific fine-tuning but is stable to arbitrary perturbations. Put differently, MBL is the fully interacting version of localisation, while Anderson localisation describes a non-interacting single-particle phenomenon.

Having said this, there is a cartoon limit of many-body localisation which is extremely transparent. We illustrate this using a disordered transverse field

Box 10.1: Phase structure in and out of equilibrium

Ising model in $d = 1$:

$$-H_{\text{TFIM}} = \sum_{j=1}^{N} J_j \sigma_j^z \sigma_{j+1}^z + \sum_{j=1}^{N} \Gamma_j \sigma_j^x \ , \tag{10.4}$$

where exchange couplings J_j and fields h_j are random variables, and we use periodic boundary conditions $\sigma_1 = \sigma_{N+1}$.

The case with $J_i \equiv 0$ is very simple–each spin aligns with its local field h_i, and the eigenstates are the 'classical' configurations of the spins along the field axis. In this case, states nearby in energy clearly need not look similar: if one finds a large set of sites λ with $\sum_{i \in \lambda} h_i \approx 0$, one can flip this entire set and end up with a very different, but near-degenerate, state: eigenstate thermalisation is manifestly violated.

The amazing feature of MBL, despite all its sublteties, is that this picture is a good starting point for the description of the generic situation. This idea is captured by the idea of so-called l-bits, which states that a *local* change of variables turns an MBL-Hamiltonian into an entirely classical-looking *local* one. That is to say, one can define a set of l-bits, τ_j^z, to diagonalise the Hamiltonian as follows:

$$\tau_j^z = \sum_{k,\alpha} A_{j,k}^{\alpha} \sigma_k^{\alpha} + \sum_{kl,\alpha\beta} A_{j,kl}^{\alpha\beta} \sigma_k^{\alpha} \sigma_l^{\beta} + \sum_{klm,\alpha\beta\gamma} A_{j,klm}^{\alpha\beta\gamma} \sigma_k^{\alpha} \sigma_l^{\beta} \sigma_m^{\gamma} + \dots \tag{10.5}$$

$$H = \sum_j h_j \tau_j^z + \sum_{jk} h_{jk} \tau_j^z \tau_k^z + \sum_{jkl} h_{jkl} \tau_j^z \tau_k^z \tau_l^z + \dots \tag{10.6}$$

Crucially, the coefficents A and h in these expressions vanish rapidly unless their indices refer to sites nearby to the reference site j. The above cartoon essentially consists of keeping only the first term in each line (and more formally corresponds to a non-interacting limit in which the change of basis between physical and l-bits is linear).

10.2.6 Eigenstate (or eigenspectrum) order

The Hamiltonian in Eq. 10.6 being local and classical (i.e., consisting of fully mutually commuting terms), the spectrum does not exhibit level repulsion like in random matrix theory. Therefore, unlike the case of a system obeying eigenstate thermalisation, the level statistics will be Poissonian.

There can, in addition, be further correlations between energy levels, which signal the existence of different phases even in this non-equilbrium setting. This possibility of defining crisply separated phases outside of thermodynamic

Box 10.1: Phase structure in and out of equilibrium

equilibrium is all the more remarkable for the twin facts that (i) it can exist in settings where the related disorder-free Hamiltonian obeying ETH does not support non-trivial order and (ii) new types of order appear, without a counterpart in equilibrium systems. The latter is the subject of the Section 10.6 on time crystals, and the former the subject of the following discussion.

The above cartoon argument could alternatively been made setting not the $J_i \equiv 0$, but the $h_i \equiv 0$. Not much would have changed. On the surface, the natural basis choice would have been the classical one for the exchange, and more deeply, the Ising symmetry would then result in *pairs* of degenerate states, related by a global Ising spin flip.

It turns out that this feature also persists beyond the cartoon limit: the full many-body states remain paired into quasidegenerate doublets. Such a degeneracy in turn implies that in the Hamiltonian, Eq. 10.6, the 'Ising-odd' terms, namely those with an odd number of τ^z's must vanish: $h_j = h_{jkl} = 0$, while $h_{jk} \neq 0$.

Note that a spectral degeneracy has already appeared in Section 5.1.4, where a three-fold ground state degeneracy even in the absence of a local order parameter acted as a diagnostic of topological order in the fractional quantum Hall effect. The present notion of eigenspectrum order works to diagnose both topological and local types of order. It is in fact used as a standard diagnostic for the breaking of discrete symmetries in exact diagonalisation studies of finite-size lattice systems. In this setting, the ground states are Schrödinger cat states. For the transverse field Ising ferromagnets, say, these are not the oppositely magnetised ordered 'up' and 'down' states, but rather their symmetric and antisymmetric combinations,

$$|\pm\rangle = \frac{1}{\sqrt{2}} \left[|\text{up}\rangle \pm |\text{down}\rangle \right] \ . \tag{10.7}$$

This mixing results from the possibility of sweeping a domain wall across the system to connect the two states; as a state with a domain wall has a nonzero activation energy, this requires a virtual process of $O(N)$ steps, so that the resulting splitting is exponentially small in N, and vanishes in the thermodynamic limit $N \to \infty$.

Conventionally, this degeneracy is present in the ground state only. Excited states, by contrast, tend to form bands. For instance, in the Ising chain, all states with a single domain wall are degenerate as far as the exchange is concerned, while the transverse field allows the domain wall to move by flipping

---

Box 10.1: Phase structure in and out of equilibrium

---

a spin adjacent to it:

$$\langle\uparrow\uparrow\uparrow\downarrow\downarrow\,|(-\Gamma\sigma^x)|\,\uparrow\uparrow\downarrow\downarrow\downarrow\rangle = -\Gamma \ . \qquad (10.8)$$

This thus gives rise to a hopping problem with dispersion $-2\Gamma\cos(k\,a)$ with $k = 2\pi\,n/N$ so that so that the resulting splitting is only algebraically small in N.

However, for nonuniform values of J_j, the disorder in J_j localises the domain walls. This happens because the energy of a domain wall depends on its location: the process depicted in Eq. 10.8 no longer connects states at the same energy. Therefore, even excited states containing such domain walls need not mix to yield the momentum eigenstates labelled by k that appear int eh dispersive band for the clean system.

As a result, eigenspectrum order indicating the breaking of an Ising symmetry can be present even at finite energy densities above the ground state, where ergodic systems in low dimension cannot support such order. Figure 10.3 illustrates the notion of eigenspectrum order for the case of periodically driven Ising chains.

The use of eigenspectrum order at finite energy densities as an ordering diagnostic is limited by the fact that the energy level spacing itself is exponentially small, $O(2^{-N})$, so that an exponentially nearby second cat state can get lost in the sea of other nearby states. However, both cat states exhibit the same long-range order in their spin correlations. In disordered magnets, such order is not measured by a straightforward Fourier component of the magnetisation (such as the one at $k = 0$ for a ferromagnet and at $k = \pi$ for an antiferromagnet), but by an Edwards-Anderson order parameter, Eq. 10.18.

---

[a] The precise definition of what constitutes a relevant conserved quantity is not entirely settled. Indeed, the projection $|a_i|^2$ of an arbitrary wavefunction, $|\psi(t)\rangle = \sum_i a_i(t)|\phi_i\rangle$ on any given eigenstate, $|\phi_i\rangle$, of the time evolution operator, is time-independent by the very definition of an eigenstate. The number of such constants of motion equals the size of the Hilbert space, and is hence exponentially large in system size for, say, a spin system on a lattice. However, a projection onto an individual eigenstate, $|\phi_i\rangle\langle\phi_i|$, is in general a highly non-local operator and thus not relevant for consideration as an observable in the conventional sense.

10.3 Floquet topological insulators

We start off with a discussion which is quite analogous to the study of topological band structures presented earlier, see e.g. Section 3.2. We are therefore interested in the properties of the spectrum of the single particle states of a non-interacting Floquet Hamiltonian. As this treats single-particle states as essentially independent, it manifestly falls under the heading of non-interacting 'integrable' systems.

The central result is the following: by subjecting a non-topological static band

structure to a periodic modulation with zero average, Oka and Aoki (2009) showed that one can obtain a topological Floquet band structure, the first instance of what is now known as a Floquet topological insulator (Lindner et al., 2011).

10.3.1 Floquet engineering

This is an instance of an application of a general set of ideas which go under the heading of Floquet engineering (Oka and Kitamura, 2019; Rudner and Lindner, 2020). This appeals to a separation of timescales between a fast driving and a much longer timescale on which the response of the system is probed. Historically, the phenomenon of dynamic localisation was a first application of this kind of Floquet engineering (Dunlap and Kenkre, 1986). This was concerned with a particle hopping on a one-dimensional lattice, subject to a uniform electric field applied at frequency Ω. In the static case, $\Omega = 0$, localisation arises via Bloch oscillations. For sinusoidal driving of the potential difference between adjacent sites, $V(t) = V \sin(\Omega t)$, the particle also executes a micromotion but for fine-tuned values of V/Ω (roots of the Bessel function $J_0(V/\Omega) = 0$, it unfailingly returns periodically to its original location: when observed on a timescale much larger than the driving period, the particle appears localised, no matter how strong its bare hopping matrix element.

The basic ingredient for analysing such high-frequency Floquet engineering is a perturbative expansion which is controlled in the smallness of the driving period, $T \sim 1/\Omega$. This so-called Magnus expansion for the Floquet Hamiltonian, $H_F = \sum_i H_i$, consists of a sequence of nested commutators of depth i, the first two terms of which are simple:

$$H_0 = \frac{1}{T} \int_0^T dt \, H(t) \tag{10.9}$$

$$H_1 = \frac{1}{2} \left(\frac{1}{T} \right)^2 \int_0^T dt \int_0^{t_1} dt_1 \, [H(t), H(t_1)] \,. \tag{10.10}$$

H_0 is just the average Hamiltonian, and H_1 encodes to what degree the instantaneous Hamiltonians at different times fail to commute.

The generation of a topological band structure was first demonstrated for the case of graphene subjected to a time-varying field (Oka and Aoki, 2009). The simplest way to derive this result is to consider the following cartoon of a circularly polarised light-field (Rudner and Lindner, 2020): the drive period T is subdivided into four portions of equal duration, during which the field successively points in the x, y, $-x$, and $-y$-directions, i.e. it rotates by 90^o at each step. This can be encoded by a vector potential $\mathbf{A}_0^{(n)}$, $n = 1 \dots 4$, which points along these directions during the corresponding parts of the drive. The Floquet unitary over a full period for the

mode at momentum $\mathbf{k}$ is then given by

$$\mathcal{U}_F(\mathbf{k}) = \mathcal{U}^{(4)}(\mathbf{k})\mathcal{U}^{(3)}(\mathbf{k})\mathcal{U}^{(2)}(\mathbf{k})\mathcal{U}^{(1)}(\mathbf{k}) \tag{10.11}$$

$$\mathcal{U}^{(n)}(\mathbf{k}) = \exp\left(\frac{iH^{(n)}(\mathbf{k})T}{4\hbar}\right) \tag{10.12}$$

$$H^{(n)} = v_F\left(\hbar\mathbf{k} - e\mathbf{A}_0^{(n)}\right)\cdot\sigma\ , \tag{10.13}$$

where the Pauli matrix refers to the graphene sublattice as in Section 2.5.

At the Dirac point, the system is gapless in the absence of driving. As the fast drive is switched on, this ceases to be the case, as can be seen by expanding the above equation to second order in the small parameter $\kappa = ev_F|\mathbf{A}_0|T/(4\hbar)$ to obtain a mass term:

$$\mathcal{U}_F(\mathbf{k}=0) = \exp(-i\kappa\sigma_y)\exp(-i\kappa\sigma_x)\exp(i\kappa\sigma_y)\exp(i\kappa\sigma_y) \tag{10.14}$$

$$\approx \mathbf{1} + \kappa^2[\sigma_x,\sigma_y] \approx \exp(2i\kappa^2\sigma_z)\ . \tag{10.15}$$

This amounts to a Floquet engineered Hamiltonian

$$H_F(\mathbf{k}=0) = \frac{2\kappa^2\hbar}{T}\sigma_z\ . \tag{10.16}$$

The field-induced mass term/gap thus scales with the intensity of the periodic electric field, and vanishes with the driving period $\kappa^2/T \sim A_0^2 T$. The sign of the mass term is set by the sign of the polarisation (clockwise or counterclockwise) of the oscillating electric field. Under time reversal, these are interchanged, and hence the the circularly polarised electric field breaks time-reversal invariance, thus removing the symmetry which protects the gapless Dirac points. In this sense, the appearance of a gap in this setting is unavoidable.

So far, so good. However, there are two flies in the ointment. Firstly, in condensed matter physics, effective Hamiltonians tend to be effective ones, in the sense of describing the low-energy physics. This limits the scope for having a simple high-frequency expansion, as there will be higher-energy degrees of freedom to which one can potentially couple.

This is related to the question of single-particle versus many-body physics: the spectrum of a many-body system is unbounded above, although the matrix elements to the very high energy states may be very small. The fundamental difference between static and Floquet systems in this language is that the non-ineracting band picture can be a good starting point even for the interacting many-body system, as the band gap between the highest filled and the lowest empty band can allow for an adiabatic switching on of the interactions. In Floquet systems, where all many-body states are crowded into the Floquet Brioullin zone of size $2\pi\hbar/T$, such a gap is generally not available.

In an actual Floquet experiment, it is hence a question of detail what level of heating one can live with. In practise, this will also depend on the observable in

question. (For a general discussion of shaking in many-body physics/optical lattices, see (Eckardt, 2017).)

Secondly, and relatedly, fixing the initial state is another challenge. Ideally in the static situation, in order to measure the Chern number of a band, one puts the chemical potential in the gap above that band, so that all the states in the band in question are filled, and empty in the bands above, as described above; and then one measures the relevant transport coefficient.

This cannot be done in Floquet systems, for essentially the same reason as described above: there is strictly no concept of high and low energy; put differently, bands do not naturally fill up one after the other as a chemical potential is increased. One thus needs to devise a separate part of the experimental protocol about to prepare an appropriate initial state.

Nonetheless, the situation is far from hopeless. For one thing, heating may be so slow to allow for a perthermal regime which persists on a timescale parametrically large in drive parameters (Abanin et al., 2015; Kuwahara et al., 2016). On this timescale, the system does not maximise entropy, and interesting phenomena may be observable starting from an appropriate, e.g. thermal, initial state. Thus, in a system driven periodically for a finite period rather than an infinite period of time, the energy spectrum is in practise not perfectly periodic but one can still see a significant number of Floquet replicas, as in the angle-resolved photoemission spectroscopy of an optically pumped topological insulator surface (Wang et al., 2013).

10.4 Anomalous Floquet-Anderson insulator

The Floquet topological insulator demonstrates how one can change the topology of a band structure using periodic driving. The resulting graphene Floquet bandstructure discussed above is essentially identical to what could have been achieved – at least theoretically – by directly adding a *static* mass term $\propto \sigma^z$. The question which naturally poses itself is whether the Floquet setting allows for topological band structures without a static counterpart. In the following, we present two such instances, both of which turn out to have rather simple rationalisations.

The first appears already for a two-dimensional band structure consisting of only two bands in which disorder can localise every single particle state, but which nonetheless exhibits stable chiral edge states; this is known as the anomalous Floquet Anderson insulator (AFAI) (Titum et al., 2016). The second is a band structure with a topologically protected Majorana zero mode away from (quasi-)energy 0, which is known as the π-Majorana fermion. In both cases, the fact that the quasienergy is periodic, i.e. that it 'lives on a circle', is the enabling new ingredient.

To see how the AFAI comes about, it is useful to recall once more Laughlin's flux insertion argument linking the existence of the gapless chiral edge mode to a bulk property, the existence of a delocalised state, and to see how its strictures

can be avoided in a Floquet system. To recapitulate this argument for the static case, consider an annulus (a circular disk with a circular hole at its origin). Upon adiabatic insertion of a *unit* flux, the initial and final Hamiltonians, H_i and H_f, are the same, up to a gauge transformation which removes the flux: a unit flux cannot yield a non-trivial Aharanov-Bohm phase.

This does, however, not mean that the state of the system remains unchanged: it is possible that the final state, $|\psi_f\rangle$, is distinct from the initial state, $|\psi_i\rangle$. This is the case for the topologically protected chiral edge state, which increases in energy ('an excitation is added' to one edge upon flux insertion). A little bit more specifically, if there is a chiral edge state connecting the lower to the upper band, flux insertion amounts to an uphill spectral flow in energy. To keep the overall spectrum unchanged, there must therefore also be a downhill flow ('an excitation is removed' from the opposite edge). If all the bulk states are localised, there is no way of transporting the charge from one edge to the other, and hence a bulk delocalised state must exist in this setting.

What changes in the case of the AFAI is that the overall spectrum need not be unchanged: due to the periodicity of the quasienergy, one only needs to demand that the sum of quasienergies remains unchanged *modulo* $2\pi/T$. That means that the spectral flow of the edge state can wrap around the periodic quasienergy direction while the bulk states all remain fully localised throughout.

A process leading to such a Floquet Hamiltonian is readily sketched. As in the case of the Floquet topological insulator, we subdivide the drive period T into several (in this case, five) segments, during which the Hamiltonian H_n is constant.

As a brief aside, we note that this kind of construction has a number of desirable features. Firstly, the piecewise constant segments can often straightforwardly be chosen to be easily visualised; by contrast, a continuously time-dependent Hamiltonian will generally have intermediate forms during its smooth development which may be rather more complex. In the study of time crystals below, we will encounter the case of alternating Hamiltonians both of which are 'classical', in the sense of consistent of manifestly commuting operators, even though the Hamiltonians do not commute with each other. Also, this structure is much more easily simulated: multiplication of a handful of unitaries is much simpler than integration over a continuous family. Nonetheless, these formulations are not just equivalent: a discontinous drive profile contains more higher Fourier components than a smooth one, so that the equilibration properties of the different drives may differ considerably.

In practise, one thus devises a simple model drive the desired properties of which are present on a certain level of intuitive obviousness. This then needs to be supplemented by a detailed study of the robustness of the desired phenomenon. This is typically done numerically, and we encourage the reader to follow up on this important, but not very pedagogically instructive, aspect in the original literature.

Returning to teh AFAI drive protocol, this consists of one piece where disorder generates Anderson localisation, Fig. 10.1(a); and a second in which a four-step
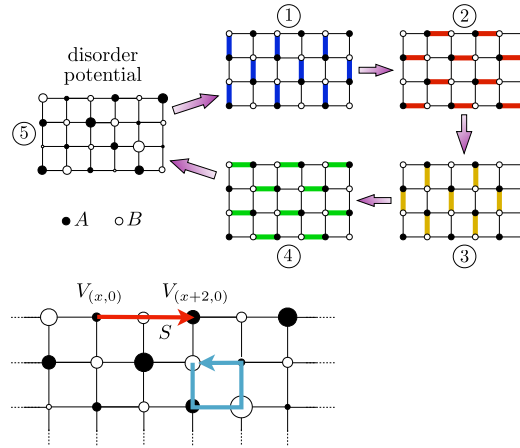
Figure 10.1 An anomalous Floquet Anderson insulator can be obtained via the quinary drive sketched in the top panel. The crux is the sequence of hoppings, labelled 1-4, which leads to a particle simply moving around a plaquette in the course of the period, as indicated in the bottom panel. (The disorder term, 5, in the drive is responsible for localising the states in the bulk.) At an edge, however, the motion is systematically disrupted in favour of a clockwise displacement. From (Titum et al., 2016).

hopping Hamiltonian provides for a motion with chiral edge states. The latter is constructed using the fact that the sites of the square lattice can be subdivided into two sublattices, A and B, so that bonds only join different sublattices; and that the bonds can be subdivided into four groups, those pointing north/east/south/west from sublattice A to sublattice B. A particle in the bulk hopping along the bonds in the driving sequence executes a clockwise motion around a plaquette. However, at the boundary, the particles on one sublattice have to skip the hop corresponding to the missing bond at the boundary, therby missing a sublattice change. The hopping sequence is constructed so that then continuing its motion from the other sublattice just moves it along the edge as shown in Fig. 10.1. This skipping motion leads to the desired delocalised edge state.

It is possible to define a bulk topological invariant which goes along with the edge states discussed here. It is the time coordinate within a period which furnishes the extra dimension entering the definition of the invariant compared to the static case; the original paper on bulk-edge correspondence in Floquet systems by Rudner et al. (2013) contains a detailed account of this construction. A more recent review of the universe of Floquet drives is (Harper et al., 2020).

10.5 Driven Kitaev chain and π-Majorana Fermions

In the anomalous Floquet-Anderson insulator, the periodicity in quasienergy of the Floquet Brioullin zone played an important role, as it allowed for the existence of edge states winding around that periodic direction.

This periodicity underpins the existence of another type of protected feature of the band structure, the so-calle π-Majorana fermion. This can be accounted for straightforwardly following the line of argument in Section 9.4 for the robustness of the Majorana zero mode, which to preclude confusion we will refer to as 0-Majorana in the following.

The crucial item of the argument concerned the symmetry ensuring that states come in pairs at energies $\pm\epsilon$, so that an isolated state in a gap would be pinned at $\epsilon = 0$. This generalises to the quasienergy in an obvious way for the state at $\epsilon = 0$, so that the (quasienergy-)0 Majorana fermion is directly corresponds to the (energy-)0 Majorana mode in the undriven case.

Now, the quasienergy being defined modulo $2\pi/T$ means that there is another location which is special: $\pi/T = -\pi/T$ modulo $2\pi/T$. If there is a gap at π/T, an isolated mode at this quasienergy will also be pinned at that energy, and thus analogously topologically protected. Such a mode is known as the π-Majorana. Crucially, this has no undriven counterpart: the periodicity of the quasienergy cannot vanish continuously, as the continuous time-translation symmetry is either present or not. It is therefore present only in the non-equilibrium setting.

These two modes can exist entirely independently of each other, so that one can have four combinations of Majoranas: trivial (i.e., none), 0, π and 0π. These can be realised straightforwardly by constructing a binary drive along the lines described above: one combines two Hamiltonians to construct a driven version of the Kitaev chain discussed in Section 9.4.

As explained there, see also Box 9.3, this can alternatively be viewed as a (driven) Ising chain, and we adapt that picture for ease of visualisation, as the binary drive is physically particularly transparent there.

The first member of the binary drive is an Ising exchange, while the second is a transverse field. The unitary over a Floquet cycle then take the form $\mathcal{U} = \mathcal{U}_\Gamma \mathcal{U}_J$, with

$$\mathcal{U}_J = \exp\left[-i\sum_j J_j \sigma_j^z \sigma_{j+1}^z\right] \ , \ \mathcal{U}_\Gamma = \exp\left[-i\sum_j \Gamma_j \sigma_j^x\right] \ , \qquad (10.17)$$

where we will be considering both open and periodic boundary conditions in what follows. Note that we have committed an abuse of notation by suppressing the explicit role of the drive period T in Eq. 10.17, by identifying $J_j T/2$ and $\Gamma_j T/2$ with J_j and Γ_j, respectively. Like this, these two variables still encode the relative strengths of the exchange and the field; and they also signal more directly a periodicity in their strength in this setting: there is no high-field limit as such. Rather,
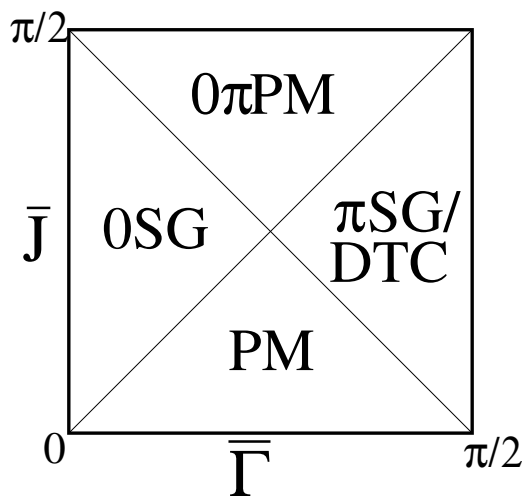
Figure 10.2 Phase diagram of the Ising spin chain subject to the binary drive, Eq. 10.17 (Khemani et al., 2016).

shifting each coupling constant by, say, 2π has no effect, so that the phase diagram in the space of couplings is itself also periodic.

Since this is a disordered problem, the distributions of J_j and Γ_j need to be specified. The result does not depend on these distributions in detail, but what does matter are the mean values $\overline{J}$ and $\overline{\Gamma}$; and for concreteness, it may be useful to think of a box-like probability distribution of width of a small fraction of 2π. Fig. 10.2 shows the phase diagram for these parameters.

Let us now consider the four possible situations, with none, one (0 or π), or both (0 and π) Majoranas in turn. The conceptually most straightforward cases are those which are analogues of the static cases, namely those without the π Majorana. These can be considered by setting one of the couplings to vanish entirely, so that the resulting problem is the static one discussed in Box 10.1 on phase structure.

The eigenstates of the paramagnet then are just the 'classical' ones diagonal in the preferred basis of $\mathcal{U}_\Gamma$, see Section 10.2.5. Switching the non-commuting exchange term, $\mathcal{U}_J$, back on then leads to a dressing of these states but, like in a paramagnet, all correlators remain short-ranged.

10.5.1 0 and π spin glasses

Considering first the case of vanishing fields $\Gamma_j \equiv 0$, we end up with a disordered classical magnet, as in Section 10.2.6. As the couplings are disordered, the order parameter diagnosing the difference to a paramagnet is not a simple (anti)ferromagnetic one in the form of a magnetisation at a given wavevector. Instead, it is an Edwards-

Anderson spin glass order parameter, which is designed to distinguish whether correlations are large (but of random sign), or instead are small in modulus. This is done by simply squaring the correlator before averaging over space:

$$q_{\mathrm{EA}} = \frac{1}{N^2} \sum_{i,j=1}^{N} \langle \sigma_i^z \sigma_j^z \rangle^2 \; . \tag{10.18}$$

Notably, the Edwards-Anderson order paramter, q_{EA}, can non-zero no matter which eigenstate it is evaluated in for the fully many-body localised spectrum we consider: this is an instance of a type of order in $d = 1$ which is cannot be present in a thermal setting.

The spin glass state with non-vanishing q_{EA} is called the 0-spin glass, 0SG, to distinguish it from the πSG (Khemani et al., 2016), which we discuss next.

The case of a π-Majorana fermion is not reducible to a simple undriven system, as explained above. The cartoon picture of this state is a combination of $\mathcal{U}_J$ as in the 0SG, but with an intervening $\mathcal{U}_\Gamma$, with $\overline{\Gamma}$ chosen so that it effects a global spin flip $\mathcal{P}_\pi$ interchanging the states with the opposite Ising polarisations: $|\uparrow\downarrow\downarrow\uparrow\downarrow\rangle \leftrightarrow |\downarrow\uparrow\uparrow\downarrow\uparrow\rangle$. The overall unitary thus reads

$$\mathcal{U}_{\pi\mathrm{SG}} = \mathcal{U}_J \, \mathcal{P}_\pi \; . \tag{10.19}$$

It is easy to see that the two finite-size eigenstates, the Schrödinger cats of Section 10.2.6, remain eigenstates of this $\mathcal{U}_{\pi\mathrm{SG}}$ but – crucially – the antisymmetric one picks up an overall minus sign due to the spin flip:

$$\mathcal{U}_{\pi\mathrm{SG}} |\psi_\pm\rangle := \mathcal{U}_{\pi\mathrm{SG}} \frac{1}{\sqrt{2}} \left[|\uparrow\downarrow\downarrow\uparrow\downarrow\rangle \pm |\downarrow\uparrow\uparrow\downarrow\uparrow\rangle \right]$$

$$= \pm \exp(i\phi) \frac{1}{\sqrt{2}} \left[|\uparrow\downarrow\downarrow\uparrow\downarrow\rangle \pm |\downarrow\uparrow\uparrow\downarrow\uparrow\rangle \right] = \pm \exp(i\phi) |\psi_\pm\rangle \; , \tag{10.20}$$

where the angle ϕ is the same for both states. As advertised, the minus sign corresponds to an extra phase $-1 = \exp(i\pi)$ picked up in the time evolution over the course of the period. Therefore, the two states are now separated by a quasienergy π. This amounts to them being located a maximal distance from each other in quasienergy, see Fig. 10.3. Note that the two states $|\psi_\pm\rangle$ are locally indistinguishable, yet they are non-degenerate; this reflects the fact that the unitary $\mathcal{U}_{\pi\mathrm{SG}}$ *cannot* be written in terms of a local, static Floquet Hamiltonian. This is a reflection of the fact that the π-SG is a genuine new *Floquet* phase.

This quasienergy difference in the single-particle picture is 'supplied' by the occupancy of the π-Majorana mode: the two states in the doublets for the 0SG and those for the πSG differ by the occupancies of the 0 and π-Majorana modes, respectively.

10.5.2 The 0π paramagnet as a symmetry-protected topological phase

The situation for the joint occupancy of both 0 and π-Majorana modes then naturally gives a quartet of states, consisting of two quasidegenerate pairs a distance π apart. Starting from any given state, one can cycle between these by toggling occupancies of the Majorana modes at 0, π, and then again at 0. The phase characterised by this eigenstate order is called the 0π paramagnet. Unlike the prior two phases, it has no spin-glass order; but unlike the trivial paramagnet, it is a symmetry-protected topological state, see Chapter 11. Like in the case of the AKLT chain, Section 5.2.4, this means that its behaviour with periodic boundary conditions is trivial, but for open boundary conditions, topologically protected edge modes appear.

This can be seen by analysing the unitary underlying the 0π Floquet paramagnet (0πPM), again by considering the simple special case $J_i \equiv \pi/2$, where

$$\mathcal{U}_{0\pi\mathrm{PM}} = \prod_j \exp[i\frac{\pi}{2}\sigma_j^z\sigma_{j+1}^z] = \begin{cases} (-i)^N \text{ periodic b.c.} \\ (-i)^{N-1}\sigma_1^z\sigma_N^z \text{ open b.c.} \end{cases} \tag{10.21}$$

The exchange terms hence contributes a state-independent global phase in the case of periodic boundary conditions; the physical behaviour of the periodic chain is hence essentially that of the trivial paramagnet.

For an open chain, however, the unitary depends on the state of the spins at the endpoints of the chain. The resulting unitary is thence that of the paramagnet, multiplied by $\sigma_1^z\sigma_N^z$. Regarding all the spins in the interior of the chain, $\mathcal{U}_J$ still has no influence, and all the action takes place at the surface. The surface spins are subject to their local fields, h_1 and h_N, and the 'exchange' term $\sigma_1\sigma_N$ from $\mathcal{U}_J$. The state of the full system thence factorises into a product state of the interior and the edge spins, and one can essentially ignore the interior part. The four eigenstates of the edge spins are then labelled by two Ising variables. First, the product of the spins expressed in the basis along the field direction; and second, the parity which encodes whether the edge spins are in an even or an odd superposition with their Ising-reversed copies. One of these labels the states within the degenerate doublet, the other the doublets separated by quasienergy π, as depicted in Fig. 10.3.

Putting all of this together yields an inert bulk, while the edge spins do exhibit *period-doubling*, i.e. their dynamics has a component at a frequency equal to half of the drive's. This can most easily be seen by explicitly time-evolving the spin operators at the edge with the unitary $\mathcal{U}$, Eq. 10.21.

Such period doubling is a most remarkable phenomenon, and it is its presence in the π spin glass which is responsible for it being called a discrete time crystal, as we discuss in more detail next.
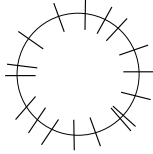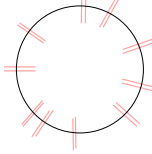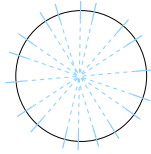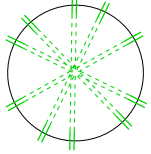
PM	0SG	πSG	0πPM

$\|\rightarrow\rightarrow\leftarrow\rightarrow\rightarrow\rangle$	$\|\uparrow\downarrow\downarrow\uparrow\downarrow\rangle$ $\|\downarrow\uparrow\uparrow\downarrow\uparrow\rangle$	$\|\uparrow\downarrow\downarrow\uparrow\downarrow\rangle$ $\|\uparrow\downarrow\downarrow\uparrow\downarrow\rangle$ $+\|\downarrow\uparrow\uparrow\downarrow\uparrow\rangle$ $-\|\downarrow\uparrow\uparrow\downarrow\uparrow\rangle$	$\|\rightarrow\rightarrow\leftarrow\rightarrow\rightarrow\rangle$ $\|\rightarrow\rightarrow\leftarrow\rightarrow\leftarrow\rangle$ $\pm\|\leftarrow\rightarrow\leftarrow\rightarrow\leftarrow\rangle$ $\pm\|\leftarrow\rightarrow\leftarrow\rightarrow\rightarrow\rangle$
$\langle\sigma^z\rangle=0$ $\langle\sigma^x\rangle\neq0$	$\langle\sigma^z\rangle\neq0$	subh. $\langle\sigma^z\rangle\neq0$ discrete time crystal	$\langle\sigma^z\rangle=0$, $\langle\sigma^x\rangle\neq0$ subh. $\langle\sigma^x_{\text{edge}}\rangle\neq0$

Figure 10.3 Observables corresponding to the phases of the Ising spin chain subject to the binary drive, Eq. 10.17. The quasienergy 'axis' is compactified into a circle of unit radius, as it is periodic with 2π. The eigenstates are distributed on this circle randomly, in pairs either exponentially close, or exponentially close to being separated by π; or in quartets consisting of two pairs separated by π.

10.5.3 Temporal correlations and the Floquet discrete time crystal

The temporal correlations of the driven Ising chains turn out to contain the fundamental novelty of these Floquet phases. Having already established long-range spatial (spin-glass) order in the πSG, we now ask about its correlations in time. To set the stage, recall that equilibration in the conventional sense implied the presence of a time-independent steady state. Here, time-independence is not a natural option as the Hamiltonian itself changes over the course of a period, and the concept of equilibration is replaced by that of synchronisation: one now observes the correlations stroboscopically, say at the beginning of each drive period, and asks if the sequence thus obtained is time-independent. We will see that it is not.

Macroscopically distinct Schrödinger cat states cannot be stabilised in experiment. It is therefore very difficult to study the long-time behaviour of these eigenstates of $\mathcal{U}_{\pi\text{SG}}$. Instead, a natural starting state is the simple state

$$|\uparrow\downarrow\downarrow\uparrow\downarrow\rangle = \frac{1}{\sqrt{2}}\left[|\psi_+\rangle + |\psi_-\rangle\right] .$$

After one period, this becomes (omitting the global phase $\exp(i\epsilon)$ due to the quasienergy of $|\psi_+\rangle$):

$$\frac{1}{\sqrt{2}}\left[|\psi_+\rangle - |\psi_-\rangle\right] = |\downarrow\uparrow\uparrow\downarrow\uparrow\rangle.$$

Another period later, the state is then back to the original $|\uparrow\downarrow\downarrow\uparrow\downarrow\rangle$.

Therefore, there is temporal symmetry-breaking as well: the periodicity of the

correlations is double the period of the drive. In analogy to (space) crystals, which have a lowered space-translational symmetry, this system is hence called a time crystal. Since it is the discrete time translational invariance of the Floquet drive which has been broken, rather than the continuous spatial translation symmetry of free space, the more precise moniker is Floquet discrete time crystal (DTC).

10.6 Many-body Floquet discrete time crystal

At this stage, the reader may feel somewhat underwhelmed. After centuries of unsuccessfully searching for a system spontaneously breaking time-translation symmetry – for the longest time, these went under the name perpetuum mobile – we have found a magnet which, essentially, takes two half-rotations in order to execute a full one. And indeed, as such, this is barely front-page news. What is remarkable is the stability of this phenomenon to all sorts of perturbations. Indeed, time-translational symmetry-breaking is a conceptually subtle business, and we refer the reader for details to the review (Khemani et al., 2019), on which parts of this chapter are based, for details.

The stability of the above cartoon pictures to tuning the drive parameters away from the simplest values 0 and $\pi/2$ follows by the same arguments as those presented for the Kiteav Majorana chain. While, e.g., the edge states of the 0πPM will leak into the bulk, they will remain localised there. And, more importantly, the quasienergy difference π is pinned to this value, so that the period doubled response will remain, even if the Ising spin flip is not a perfect inversion, $\mathcal{P}_\pi$, but offset by a small but finite angle, $\mathcal{P}_{\pi-\epsilon}$. This is in stark contrast to non-interacting spins, where an offset of the flip by an angle ϵ leads to a continuous drift in the response frequency, see Fig. 10.4.

Absolute stability

These insights, however, pertain to the single-particle picture made possible thanks to the integrability of the driven Ising chain, which it inherits from the static transverse field Ising model. What happens if generic (small) terms are added to the Hamiltonian which spoil the integrability? This question is important, as the definition of a phase requires stability to perturbations–it is phase transitions, not phases, which require fine-tuning microscopic parameters.

The answer is that the πSG persists, and indeed, that it is more stable than even conventional static spin glass phases (Else et al., 2016; von Keyserlingk et al., 2016). A flavour for why this may be the case can already be obtained by asking what is the consequence of applying an infinitesimal Ising symmetry-breaking field in the direction of the ordered moments. In the case of the 0SG, this immediately breaks the exponentially small quasidegeneracy, as a state $|\uparrow\downarrow\downarrow\uparrow\downarrow\rangle$ will in general have a non-vanishing magnetisation. By contrast, the above cat states

$|\psi_\pm\rangle = \frac{1}{\sqrt{2}}\left[|\uparrow\downarrow\downarrow\uparrow\downarrow\rangle \pm |\downarrow\uparrow\uparrow\downarrow\uparrow\rangle\right]$ have a splitting which is π, i.e. $O(N^0)$; they will therefore not mix appreciably when an infinitesimal field is applied. This result can be reexpressed in a real-time picture. Formally, instead of inverting the spins by $\mathcal{U}_\Gamma$, one can choose to enter a moving reference frame which sees an inversion of the 'Ising-odd' applied field instead, while leaving the 'Ising-even' exchange invariant. This means that the field points alternatingly up or down. It hence averages out, and with the field, the Ising symmetry breaking disappears.

The reader who is uncomfortable with such wordy explanations will need to delve somewhat more deeply into the physics of time-translation symmetry-breaking, MBL, and Floquet unitaries, see e.g. Section 5 of (Khemani et al., 2019). Put concisely, the 'topological' formulation of the difference between 0 and πSG lies in the action of the unitary on the $l-$bits τ^z: $\mathcal{U}^\dagger \tau^z \mathcal{U} = \pm 1$. It can be verified (e.g. for the special point of the perfect spin flip) that the case $+1$ corresponds to the 0SG, and -1 to the πSG. These two cases cannot be smoothly be deformed into one another. Now, when an Ising symmetry-breaking field is applied, the quasi-degeneracy of the doublet of the 0SG is lifted, with the cat states $|\psi_\pm\rangle$ being replaced by the conventional, unentangled states $|\uparrow\downarrow\downarrow\uparrow\downarrow\rangle$ and $|\downarrow\uparrow\uparrow\downarrow\uparrow\rangle$ as eigenstates. Crucially, these are no longer eigenstates of $\mathcal{P}$, which therefore can no longer be used as a generator of an Ising symmetry. By contrast, there is no degenerate perturbation theory to be done in the case of the πSG, and the cat states remain eigenstates.

The second crucial ingredient to establish the absolute stability of the πSG is supplied by many-body localisation, which supplies the fact that the flip operator $\mathcal{P}$ can be perturbatively continued to $\tilde{\mathcal{P}}$ as *any* local symmetry breaking term is added to the Hamiltonian. This proceeds the same way that the l-bits evolve upon addition of those terms. Like for the l-bit, the outcome is non-universal in the sense that the resulting operator depends explicitly on the Hamiltonian, unlike the spin flip operator $\mathcal{P}_\pi$ which implements a global Ising symmetry. In this sense, $\tilde{\mathcal{P}}$ in general reflects an *emergent* Ising symmetry, i.e. one not present in the starting Hamiltonian.

It is worth noting how far outside the familiar realm of statistical physics this result lies: Landau-Ginzburg theory states from the very outset that symmetries which are not there cannot be broken. Indeed, the very notion may seem as absurd as the persistence of the smile of the Cheshire cat, even once the cat has vanished. As Alice mused, "I've often seen a cat without a grin, but a grin without a cat! It's the most curious thing I ever saw in my life!" (Carroll, 1865).

This feature, termed absolute stability, is rather reminiscent of the stability of the topological order in the fractional quantum Hall effect, Section 5.1.4, and the RVB liquid or Kitaev's toric code, which are also stable to an arbitrary small perturbation. Not unsurprisingly, it is topology which underlies both phenomena.

In all cases, one finds an emergent discrete degeneracy, involving states threaded by fluxes in the above examples of topological order. For the DTC, the Ising symmetry was the crutch with which to discover this emergent Ising symmetry, which

however can stand on its own. What basically happens is that the physics of MBL ensures that any variation in the properties of the system induced by the perturbation takes place smoothly. This is the essence of the *l*-bits, which get dressed continuously. As long as there is a discrete (i.e., topological) distinction between the objects under consideration, such as a global even or odd parity index, this distinction will persist even as everything else changes smoothly. In this way, the πSG is special. The PM is trivial to start with; the 0SG is not protected on account of its quasi-degeneracy, and the parity of the 0π paramagnet pertains to the isolated edges which are sensitive to a local symmetry-breaking perturbation, as may be checked by explicitly solving for the eigenstates of the two spins at the edges under application of a symmetry-breaking field. The formal demonstration of these properties proceeds by constructing the relevant perturbation theory and examining its impact on the eigenstate-ordered multiplets (Khemani et al., 2019).

We note, however, that even absolute stability has its limits. The one symmetry one cannot discard is the discrete time translational symmetry of the Floquet drive. Without it, there will be Fourier components to the drive which can generate matrix elements between states differing by quasienergy π. This leads to the same situation as the static field applied to the 0SG, which again proves fatal to the eigenstate order, thereby merging the πSG with the trivial PM. A Cheshire cat without a smile will not leave one behind.

We do remark on the similarity of this topology-induced stability to the case of the quantum Hall effect. Our present treatment started by providing a cartoon picture of the DTC, where the period doubling seemed to be put in entirely by hand, by splitting a full rotation into two pieces. This is just like the case of the quantum Hall effect, where it is known that a clean system always exhibits a Hall conductivity $\sigma_{xy} = \nu e^2/h$, where ν is an integer at integer filling. The amazing topological stability of the Hall effect appears when ν is tuned away from an integer or rational fraction, ν_0, but the Hall conductivity remains pinned at $\sigma_{xy} = \nu_0 e^2/h$. In the same way, in the DTC the period doubling remains robust even if the rotation angle is changed continuously away from half a rotation. While the quantum Hall effect enables one to determine the conductance quantum e^2/h to unprecedented accuracy, in the case of the discrete time crystal, the corresponding stability is that of (the admittedly pre-quantised) integers.

Spatiotemporal order

We now return to the phenomenology of the DTC, in particular to the form of the correlations it encodes: there are long-range correlations both in space and in time – it represents an entirely new form of *spatiotemporal* order. This is probed most crisply by taking the limit of large distances and long times simultaneously, where the only signal in the spin correlations is the period doubled one. By contrast, as outlined above, the long-range spin glass order corresponds to the limit of large distances at equal times. Similarly, if one measures the spin autocorrelation function,
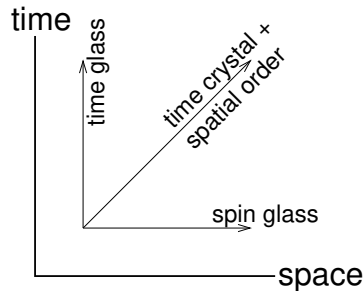
Figure 10.4 Sketch of the spatiotemporal nature of correlations in the Floquet DTC. At long distances and long times, a novel combined spatiotemporal from of order is observed. Moving along the time- (or the space-) axes only leads to 'local' glassy correlations. In particular, on-site correlations at long times exhibit oscillations at many frequencies, characteristic of the local disordered environment. Equal-time correlations at large distance evidence a standard spin glass.

i.e. the limit of long times at short distances, one is sensitive to the oscillations of the spins on account of the difference between physical variables, σ^α, and the l-bits, so that one will observe Rabi-type oscillations set by the local environment (e.g. the effective field experienced by the l-bits). This behaviour is called a time glass. This is illustrated in Fig. 10.4.

Several experiments have been undertaken to look for discrete time crystals, on platforms as varied as nitrogen-vacancy centres in diamond, in various NMR platforms, and in a chain of trapped ions; a recent review is (Khemani et al., 2019; Else et al., 2020). These also tend to use stepwise drives for clarity and ease of implementation. The central ingredients for the demonstration of a Floquet many-body DTC are the following. First, a spin rotation by an angle close to, but not at, half a rotation. Second, a dose of disorder to induce MBL; and third, interactions to lock the response robustly into period doubling.

The various experiments have seen such a frequency locking, which is a remarkable feat, and the fact that behaviour of this type has been seen across platforms is a great experimental achievement for each of these. This is all the more remarkable as the experiments arrived only a short period after the theoretical work. Indeed, they have clarified considerably the understanding of discrete time crystals. In particular, a detailed analysis of these experiments show that none of them has as yet realised the Floquet DTC *sensu stricto*. However, while not yet *bona fide* non-equilibrium phases of matter, they have unearthed a set of very interesting related phenomena, which now go under headings like symmetry-protected, algebraic or prethermal time crystals. The latest proposal in this realm is to use noisy intermediate-scale quantum (NISQ) technology to realise the Floquet DTC, as such a platform is ideally suited for the emulation of a stepwise drive, with locally addressable quantum gates naturally implementing the disordered pairwise interactions and fields appearing in
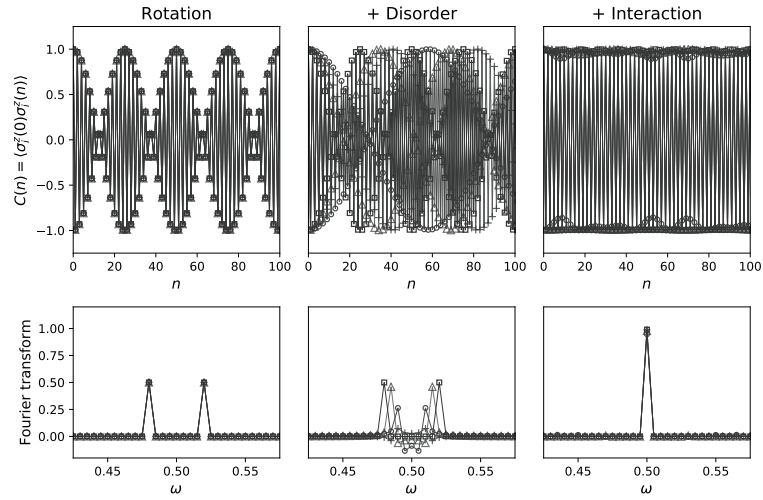
Figure 10.5 Signatures of Floquet many-body discrete time crystal. The drive consists of an imperfect half-rotation (left panel), with added disorder in the field (middle panel). Finally, upon addition of a disordered exchange, the spins exhibit robust period doubling. The top panels show the corresponding stroboscopic real-time evolution of the temporal on-site correlators of the σ^z. The bottom panels show the same information in Fourier space, where the locking into period doubling manifests itself in a peak locked at $\omega/\Omega = 1/2$. Figure courtesy of Matteo Ippoliti.

the model Hamiltonian, and local initialisation and read-out enabling a detailed analysis of the DTC signatures in the correlations for arbitrary initial states.

The theorist's dream experiment is shown in Fig. 10.5. The data there shows the polarisation of four different spins in a disordered interacting spin chain of length $L = 16$ as a function of stroboscopic time. This means that the value of the spin is obtained at a sequence of times offset by an integer number of drive periods $T = 2\pi/\Omega$. It is obtained from numerical simulations as follows.

The leftmost panel shows the effect of a uniform field effecting not quite half a rotation. This leads to peaks in the Fourier transform of the correlation function away from π. The second part supplies a disordered field which puts the individual spins out of step and leads to differences between the peaks, which still do not exhibit precise period doubling. The third part of the drive adds a disordered Ising exchange. It is only when this is added that the response locks in at period doubling, and the peak robustly shifts to frequency π. Note that the starting state was a random product state: the DTC signal is present for all initial states.