REVIEW

Quantum spin liquids

Tom Lancaster^a

^aDepartment of Physics, Durham University, South Road, Durham, DH1 3LE, UK

ARTICLE HISTORY

Compiled December 19, 2023

ABSTRACT

A glance at recent research on magnetism turns up a curious set of articles discussing, or claiming evidence for, a state of matter called a quantum spin liquid (QSL). These articles are notable in their invocation of exotic notions of topological physics, quantum entanglement, fractional quantum numbers, anyon statistics and gauge field theories. So what is a QSL and why do we need this complicated technical vocabulary to describe it? Our aim in this article is to introduce some of these concepts and provide a discussion of what a QSL is, where it might occur in Nature and why it is of interest. As we'll see, this is a rich subject which is still in development, and unambiguous evidence for the realisation of the QSL state in a magnetic material remains hotly debated. However, the payoff in terms of the special nature of quantum entanglement in the QSL, and its diverse spectrum of unusual excitations and topological status will (at least to some extent) justify the need to engage with some powerful, occasionally abstract, technical material.

KEYWORDS

quantum spin liquid, gauge theory, topology, spinon

1. Introduction

We start with an attempt at a working definition of a quantum spin liquid (QSL) taken from Ref. [1]: a quantum spin liquid ground state is an electronic insulator with spinrotation symmetry and an odd number of electrons per unit cell. Although there's quite a lot going on here, one key feature is that we have an interacting electronic system comprising localised electrons whose spin moments do not align into a magneticallyordered configuration, even down to T = 0. Although this notion of a lack of spin order will provide a jumping-off point, we'll see that the picture of a disordered magnetic state fails to capture much of the rich physics that underlies the QSL. However, the general principle of order will underlie our discussion. Specifically, we shall start from Lev Landau's description of the order that follows from a symmetry-breaking phase transition [2], and the subsequently development of this subject by Philip Anderson [3], since these concepts have provided a means of classifying and understanding ordered phases of condensed matter.

The discovery and elucidation of the fractional quantum Hall (FQH) fluid fundamentally challenged this pervasive notion of order [1,4]. The FQH fluid is a state of matter realised in a two-dimensional electron system in high magnetic fields [5]. The

CONTACT Tom Lancaster. Email: tom.lancaster@durham.ac.uk

fluid has a gap in energy between the ground state and the lowest-lying excited states of the system.¹ If we put enough energy into the system to excite excitations over the energy gap, the resulting particle-like states do not carry integer numbers of elementary electronic charges, but rather fractions of an electron charge. We say that they are *fractionalised*. Moreover, they are neither bosons nor fermions, but are known as anyons, and have their own, unique, quantum statistics. There is not one FQH fluid, but many, each realised in practice by applying a magnetic field of a different magnitude. Moreover, we cannot distinguish or classify the different FQH states by their symmetries, as we can with other ordered states of matter. Instead, topology, the study of the shapes of objects and spaces, is invoked, and we treat the FQH fluid as being *topologically ordered* (TO). Relative to conventionally-ordered states, TO states are typified by (i) a ground-state degeneracy that depends on the topology of the system, (ii) a gap against excitations and (iii) a high degree of quantum entanglement.

Here is where the concept of a QSL fits into condensed matter. Like the FQH fluids, there is not one QSL state, but many that can't be distinguished by symmetry. All QSLs have fractionalised excitations. Some QSLs show topological order and have gapped excitations; some do not. QSLs feature an unusually-high degree of (long-range) quantum entanglement in their ground states. We will take the view that it is the nature of the quantum entanglement that characterises the QSL and explains its properties, rather than its magnetic disorder. Following X.-G. Wen [1], we call the special entanglement quantum order as a shorthand,² because the unique pattern of entanglements characterises each QSL. Topological order is then a special case of quantum order in which all excitations have finite energy gaps between the ground state and the excited states.³

This paper is structured as follows: we start with an introduction to the treatment of order and its absence in magnets, followed in Section 3 by a survey of the conceptual ingredients required to understand this field. Two of these: anyons and gauge theory are introduced in more detail in Sections 4 and 5. We then discuss the toric code, which is the simplest model of a QSL. Fractionalised spinons are described in Section 7 as a route to finding more model QSLs, before the exactly solvable Kitaev model is introduced in Section 8. Finally we briefly assess some examples in the search for a realisation of the QSL state.

2. Background: magnetic order, disorder and the RVB

A simple picture of a magnetically disordered system is shown in Fig. 1(a). The arrows represent spins arranged on an ordered lattice. (This structural order is always assumed, the disorder refers only to the arrangement of magnetic moment directions.) The system can be understood [2,3,7] via its symmetry: the magnetisation M (that is, the average magnetic moment) is zero, since as many arrows point up as down. If we turn each of the moments through 180° then we obtain the situation shown in

 $^{^{1}}$ The existence or absence of such energy gaps will turn out to be a useful notion in classifying states of matter.

²The entanglement partially explains the Q in QSL. The other feature is that we are dealing with T = 0, where fluctuations are purely quantum mechanical rather than thermal. In contrast, a classical spin liquid has T = 0 order destroyed by T > 0 thermal fluctuations owing to high ground-state degeneracy.

³To close the loop on our technical definition above, with its specification of odd numbers of electrons, we note that the *Lieb-Schultz-Mattis* theorem [6] says that a system with a half-odd integer spin per unit cell, and without symmetry breaking order, has either (i) a unique ground state with gapless excitations (e.g. a gapless QSL) or (ii) a degenerate ground state with a gap (e.g. a topologically ordered QSL).

(a)	ŧ	ł	ŧ	¥	(b)	ł	ŧ	¥	ŧ	(c)	⋪	ŧ	ŧ	ŧ	(d)	¥	¥	ł	¥	(e) 🛦	ł	ţ	ł	ŧ
↑	ł	ł	ł	ł	ł	ŧ	ł	ł	1	↑	1	†	†	ł	¥	ł	ł	ł	ł	ł	1	ł	1	ł
¥	¥	ł	¥	ł	ŧ	ŧ	ł	ł	ł	↑	ŧ	≜	ŧ	ł	↓	ł	ł	ł	¥	ŧ	ł	ŧ	¥	ł
≜	¥	ł	ŧ	¥	ł	ŧ	ł	ł	ŧ	↑	ŧ	ŧ	ł	ŧ	¥	ł	¥	ł	¥	ł	ŧ	¥	ŧ	ł
ł	ŧ	¥	ł	ł	ŧ	ł	ŧ	ŧ	ł	ŧ	Ť	ŧ	ŧ	ŧ	¥	¥	ł	ł	ł	Ť	ł	ŧ	ł	ł

Figure 1. (a) A disordered (paramagnetic) magnetic state. This configuration has magnetisation M = 0 and so does (b), which is produced by turning all of the spins through 180° . (c) An ordered (ferro)magnetic phase with $M = M_0$. Turning the spins through 180° results in (d) which has $M = -M_0$. (e) The (classical) Néel antiferromagnet, with ordered, alternating spins. This can be decomposed into two antiparallel, interpenetrating, aligned magnetic sublattices.

Fig. 1(b), which also has M = 0. We therefore cannot tell from the magnetisation that we have transformed the system. This inability to tell that a change has been made is a symmetry. An ordered magnet is shown in Fig. 1(c). This has magnetisation $M_0 \neq 0$ and has lost its previous symmetry: turning the arrows through 180° [Fig. 1(d)] makes a measurable difference in that it reverses the magnetisation $M_0 \rightarrow -M_0$. The symmetry has been broken on magnetic ordering. In Nature, this sort of ordering is observed to take place via a magnetic phase transition at a temperature T_c .

Philip Anderson notes some key points here [3] that apply in general to brokensymmetry states. (We assume the spins can point in any direction in what follows.) (i) Once the rotational symmetry of the spin system is broken, the ordered magnetic state becomes rigid, such that it costs energy to deform the spin structure. (ii) A new sort of excitation emerges on symmetry breaking: the magnon. This can be thought of as a particle-like excitation formed from a single, smeared out, flipped spin. (For a system composed of s = 1/2 spins, a flip represents a s = 1 excitation.) Owing to the possibility of exciting magnons with arbitrarily long wavelength, for a rotationallyinvariant Hamiltonian it costs a vanishing amount of energy to create one, and so, formally, we say that the excitation is gapless. One way to think about the gaplessness is that it is "protected" by symmetry and the ordering.

A a magnet has a microscopic description using the Heisenberg model with Hamiltonian [7] $\hat{H} = -\sum_{\langle ij \rangle} J_{ij} \hat{S}_i \cdot \hat{S}_j$. We will consider only s = 1/2 spin operators \hat{S} and the sum will be taken only over nearest-neighbour spins, denoted $\langle ij \rangle$. The ferromagnet (FM) [Fig. 1(c)] is the ground state for constant exchange J > 0. However, instead of the ferromagnet, more commonly realised in materials is the antiferromagnet (AFM) (where J < 0). Here two interpenetrating magnetic sublattices break symmetry, with the spins arranged into an alternating Néel state [shown (classically) in Fig. 1(e)]. Antiferromagnets have several differences compared to ferromagnets, especially when quantum mechanics is taken into account [3,7]. The crux is that the Néel state is not an eigenstate of the Heisenberg Hamiltonian.

It's useful at this stage to take a brief step back and consider the quantummechanical fate of a model two-level system such as the benzene molecule [8]. Benzene has two energetically-degenerate configurations of alternating double and single covalent bonds, as shown in Fig. 2(a) (states we shall call ϕ_1 and ϕ_2). Neither is an eigenstate of the system, and there is a matrix element for a transition between them. It is straightforward to show [8] that the ground state is the symmetric superposition of the states $[(\phi_1 + \phi_2)/\sqrt{2}]$, and the antisymmetric combination $[(\phi_1 - \phi_2)/\sqrt{2}]$ is an excited state. The ground state superposition gives rise to the notion of a delocalised π



Figure 2. (a) Benzene has bonds that we picture in the ϕ_1 (left) or ϕ_2 (middle) states. The superposition of these leads to the π orbital picture (right). (b) A triangular arrangement of three spins with exchange interactions between them. (c,d) Two possible RVB configurations of singlet bonds linking spins on the triangular lattice.

orbital [Fig 2(a, right)] that represents the system fluctuating, or resonating, between the two alternating-bond configurations. Specifically, if we start the system in state ϕ_1 , then the probability of finding the system in this same state some time t later oscillates as a function of t, at a frequency $\omega = 2V/\hbar$, determined by the energy gap 2V between the eigenstates. These fluctuations have nothing to do with temperature, but are purely quantum-mechanical in nature. The point here is that since the initial state ϕ_1 is not an eigenstate of the system, probability sloshes around at a rate determined by the energy-level separation. If V is small, the resulting small energy-level separation means that the oscillations have correspondingly low frequency.

Since the Néel state is not an eigenstate of the AFM Hamiltonian, it was suggested that quantum fluctuations (generalising the sloshing of the time-dependence of the probability density in the previous example) would break up an antiferromagneticallyordered configuration of spins. However, the Néel state is stable. This is because, for the large systems that comprise condensed matter, there are lots of eigenstates arbitrarily close in energy, out of which a stable Néel-like wavepacket can be constructed. The closeness of eigenstates in energy means that the oscillation frequency is therefore very small. One way to think about this feature in an extended system is by analogy with picking up a rock, where we might suspect that position-momentum uncertainty would prevent a solid object from being localised [3,9]. The reason why the rock is stable is that we seek to localise it within, say, a lattice spacing a, so the uncertainty in momentum is of order \hbar/a and the uncertainty in energy for an atom of mass m is then $\Delta E = \hbar^2/2ma^2$. However, the key is that the solid is made of N atoms, rigidly stuck together, so $\Delta E = \hbar^2/2Nma^2$ and, since N is a macroscopic number, the energy uncertainty is very small. This implies that the frequency of oscillation between closely-spaced position states is vanishingly small for a macroscopic piece of matter. The analogous argument can be made in the case of an antiferromagnet (given in the slightly more complicated terms of a rigid rotor), and it turns out that the closelyspaced energy levels in a macroscopic system similarly allow a very stable wavepacket to be constructed giving a stable Néel AFM [9,10]. The smoking-gun experimental verification of its existence came from magnetic neutron diffraction where the periodicity of the magnetic sublattices can be directly seen from the measured Bragg peaks [7].

Although the AFM survives the strictures of quantum mechanics, it's possible to

postulate a state that shouldn't. 1973 saw Philip Anderson's suggestion of a resonating valence bond (RVB) state [11]. This is based on the notion of frustration, where it is impossible to satisfy all of the interactions for a particular geometry.

Example: consider a triangle decorated with Ising spins as shown in Fig. 2(b) with Hamiltonian $H = -\sum J_{ij}\sigma_i^z\sigma_j^z$, where all bonds favour antiferromagnetism $(J_{ij} = -J < 0)$. This state is frustrated as there is no single state that satisfies all of the interactions. In this case the frustration gives rise to six lowest-energy spin configurations (with energy E = -J) and two excited states (with E = 3J). By contrast, an unfrustrated model where all $J_{ij} = J > 0$ has a ferromagnetic, two-fold degenerate ground state with energy E = -3J. (The degeneracy here simply reflects the fact that all spins can align either up or down.) We see how frustration increases not only the relative energy of the ground state but, more crucially, the entropy through the large number of degenerate ground states that result. The degeneracy becomes macroscopic in the thermodynamic limit and has a destabilising effect on symmetry-breaking order.

To form Anderson's RVB state [11-14] we extend the triangle to a spin-1/2 Heisenberg model on an infinite two-dimensional triangular lattice with purely antiferromagnetic interactions on all bonds.⁴ Note first that, quantum mechanically, the ground state of a single antiferromagnetic bond is represented by an entangled spin singlet $(|\uparrow\downarrow\rangle - |\downarrow\uparrow\rangle)/\sqrt{2}$. We therefore form a typical state by decorating the lattice with singlets coupling nearest neighbours [Fig 2(c)]. We then repeat with a different choice of singlet bonds to form another state with the same energy [Fig 2(d)]. We continue this process until we have all possible coverings, (we can include coupling beyond nearest neighbours) and then form our final trial wavefunction by adding together all of these states to form the RVB state. Notice how the RVB bears a resemblance to the benzene example, albeit in an infinite system: in both we can think of the choice of bonds resonating between different degenerate configurations. Ultimately, the RVB state did not turn out to describe a s = 1/2 Heisenberg antiferromagnet on a triangular lattice.⁵ However, the RVB is our first example of a QSL. It does not break rotational symmetry since each bond represents a s = 0 spin state. Moreover, the state itself can be thought of in terms of a complicated dynamic dance of singlet bonds that represents the underlying pattern of quantum entanglement.

The notion of a dynamic pattern of quantum correlations that underlies a magnetically-disordered state of macroscopic matter will now be expanded via a consideration of some of the background concepts needed to understand the QSL picture.

3. Ingredients of a QSL

There are a number of ingredients of QSLs that we describe here and will feature in the rest of the article. The first is **reduced dimensionality**. When interactions are confined to less than three spatial dimensions (3D), fluctuations become more effective in destabilising order. In fact, the Mermin-Wagner theorem [4,7] says that for spins with a continuous degree of freedom, magnetic order will always be destabilised above T = 0. We shall mainly be concerned in this review with the case of two spatial dimensions (2D) and time [known as (2+1) dimensions]. Although the two-dimensional Heisenberg magnet will not order for T > 0, the ferromagnetic does at T = 0, and

⁴Reminder: two Heisenberg spins $|s_1s_2\rangle$ with isotropic antiferromagnetic exchange coupling between them will have a singlet ground state: $(|\uparrow\downarrow\rangle - |\downarrow\uparrow\rangle)/\sqrt{2}$, and a triplet of excited states, comprising $|\uparrow\uparrow\rangle$, $|\downarrow\downarrow\rangle$ and $(|\uparrow\downarrow\rangle + |\downarrow\uparrow\rangle)/\sqrt{2}$.

 $^{^{5}}$ It has, however, since been invoked in several contexts, not least to describe high $T_{\rm c}$ cuprate superconductors.

there is good evidence that the antiferromagnet should order at T = 0 as well [15]. We therefore need some further source of fluctuations if we want to prevent longrange magnetic order and instead promote the quantum order that gives robust (i.e. T = 0) spin-liquid behaviour. Frustration, as considered above for the triangle, is one possibility, but there are others, including higher-order interactions or disorder. (There are also QSL states that can be realised in 3D, which we return to at the end of the article.)

The second ingredient is **topology**. Topology enters many-body physics in a number of ways [1,4,16]. One is in theories: some theories make no contact with the distances and time intervals in a problem (i.e. the geometric features) and instead rely only on the underlying shape of their space. These are topological theories: an example is the Chern-Simons theory used to describe the FQH fluid, along with some QSL states. A consequence of this lack of geometry turns out to be that the Hamiltonian of a purely topological theory (such as a the Chern-Simons Hamiltonian) is H = 0, meaning all states are degenerate with zero energy. The exact size of the large resulting degeneracy depends on the topology of the underlying space (or manifold) on which we're working. The topology of the manifold is an interesting case in itself. If we impose periodic boundaries on our 2D lattice in both spatial directions then the shape of the underlying topological space is a torus. This topology leads to a robust degeneracy: it is very difficult to split the energies of the degenerate states with perturbations. Topological order (where there is a nonzero energy gap between the lowest-energy state and all of the excited states) follows from the existence of these robust topological degenerate ground states, and is sometimes described in terms of topology "protecting" the gaps that appear in topologically-ordered states. A final appearance of topology is made by excitations [4,17]. Some excitations are extended in space and cost a large amount of energy to remove. Topological excitations, often called topological defects, include the domain wall in the one-dimensional (1D) magnet, the vortex in 2D and the monopole in 3D. In each case, the large energetic penalty involved in removing a topological excitation can be traced back to their shape (e.g. for the domain wall, we would need to flip a semi-infinite number of spins to remove the wall).

Entanglement is our next ingredient. A state is entangled if it is in a quantum superposition that cannot be written as a product, even under an arbitrary (local) change of basis states. The key example to have in mind here are two spins in an entangled singlet state $(|\uparrow\downarrow\rangle - |\downarrow\uparrow\rangle)/\sqrt{2}$. As we've said, quantum order describes the pattern of entanglements in a many-body system, although it's rather hard to rigorously define a many-particle entangled state. We saw one example in the RVB of resonating singlets, we'll see another in our discussion of the toric-code model.

The number of different **elementary excitations**⁶ of the QSL are an important means of classifying the states, especially whether they are **gapped or gapless** (i.e. whether or not there is a nonzero energy gap between the lowest-energy state and first excited state). QSLs can support local particle-like excitations that can be fermionic, bosonic or anyonic, and can have fractional quantum numbers. For example, many QSLs support gapless, chargeless s = 1/2 fermionic excitations known as **spinons** [18]. In constrast, excitations can also be topological objects. An example of a (bosonic) excitation found in many QSLs is the topological **vortex**.

In order to discuss spin liquids we shall also need some ideas from quantum field theory (QFT) [4,19]. In a QFT we describe a system in terms of a field (an object in

⁶The low-energy excited states that exist close to the ground state of an interacting, many-body system can be thought of as hosting an assembly of elementary excitations. Examples of elementary excitations are phonons in a crystal, magnons in a ferromagnet, quasielectrons in a metal etc.

which we input a position and output an amplitude or quantum operator). Particlelike excitations in QFT are quantised excitations in the field. Operators in QFT act on many-particle states that describe the system. The most important state is the ground state or vacuum $|0\rangle$. Excited states, such as particle excitations, are then added to this vacuum state, using operators like \hat{c}^{\dagger} . For example, to add a *c*-particle to the vacuum we write $\hat{c}_i^{\dagger}|0\rangle = |1_i\rangle$, where $|1_i\rangle$ is a state containing one *c*-particle at position *i*. (We can also annihilate particles: $\hat{c}_i|1_i\rangle = |0\rangle$.) One important model to have in mind in our discussion is the tight-binding model, describing the energies of electrons on a lattice with a Hamiltonian

$$\hat{H} = \sum_{ij} (-t_{ij}) \hat{c}_i^{\dagger} \hat{c}_j, \qquad (1)$$

where the subscripts label lattice sites. The idea here is that the operators annihilate the electron at site j and create it at site i, causing an electron to hop between sites, making a contribution of kinetic energy of $(-t_{ij})$. This model can often be solved in momentum space to give $\hat{H} = \sum_{k} E(k) \hat{c}_{k}^{\dagger} \hat{c}_{k}$, where E(k) is the electron dispersion. We'll see that some important models of QSLs can be boiled down to variations of this basic picture.

The final, crucial, ingredient is **gauge field theory** [4]. A gauge field is one where there is some redundancy of description. This is to say that the same physical state can be described by several different configurations of a gauge field. (You can think of this as a little like a language in which a physically-equivalent scene can be described in English, French, Arabic, Hindi, Japanese etc.) The most familiar gauge field is the electromagnetic gauge field, whose components are $A^{\mu} = (V, A^x, A^y, A^z) = (V, A)$. This is often also called the electromagnetic potential. We define the electric field Ein terms of the components of the gauge field as $E = -\nabla V - \frac{\partial A}{\partial t}$, and the magnetic field B as $B = \nabla \times A$. An important property of these equations is that there is some freedom in how the components A^{μ} are specified. In fact, if we make the gauge transformations

$$V \to V - \frac{\partial \chi(t, \boldsymbol{x})}{\partial t}, \quad \boldsymbol{A} \to \boldsymbol{A} + \boldsymbol{\nabla} \chi$$
 (2)

where $\chi(t, \boldsymbol{x})$ is some arbitrary function, then the values for the \boldsymbol{E} and \boldsymbol{B} fields are unchanged.

The notions of anyons and of a gauge theory are so important that we'll devote some more discussion to them in the following sections.

4. Anyons and topology

One of the most striking properties of QSLs is that they can host excitations that are neither conventional fermions nor bosons [4,20]. In quantum mechanics in 3D, changing the labels on two identical particles results in a change of the wavefunction according to the rule $\psi(x_1, x_2) = \pm \psi(x_2, x_1)$, where the + sign applies to bosons and the - sign to fermions. For the case of 2D space, the exchange of particles needs more careful attention. We start with two identical particles at positions x_1 and x_2 and identify two distinct ways of exchanging them. The result of processes of type I is to move $x_1 \to x_1$ and $x_2 \to x_2$. Some examples are shown in Fig. 3(a) and (b). The particles end up where they were originally, although they may move around each other. The result of



Figure 3. Examples of ways of exchanging particles. (a) and (b) are type-I processes, where the particles end up at the same positions. In case (b) one particle loops around the other once in the exchanging process. (c) and (d) are type-II processes where the particles exchange positions. In case (d) one particle loops once around the other during the exchange. [Figure reproduced from Ref. [4], reprinted with permission of Oxford University Press.]

type-II processes is to move $x_1 \to x_2$ and $x_2 \to x_1$ [Fig. 3(c) and (d)]. Again particles may move around each other several times before settling at their final positions [as in Fig. 3(d)].

We ask what the relative quantum mechanical phase difference is between processes of type I and type II. The key parameter is the angle that one particle is moved around the other. Topology comes in here: given a set of particle paths in 2D space, we can smoothly distort the paths of the particles, but we may not change the number of times particles wrap around each other without introducing singularities in the paths. Processes of type I involve rotating particle 2 around particle 1 by angle $\phi = 2\pi p$, where the winding number p takes an integer value (including zero). Processes of type II involve rotations of $\phi = \pi(2p+1)$. Each value of p describes a topologically distinct process. We suppose that these topologically distinct processes make a multiplicative contribution to the wavefunction of $\Phi(\phi)$ which is pure phase. If we carry out a sequential string of these processes then we require that the angles add, whilst the wavefunctions should multiply. That is, $\Phi(\phi_1 + \phi_2) = \Phi(\phi_1)\Phi(\phi_2)$, which implies that $\Phi(\phi) = e^{i\eta\phi}$, where η is a parameter. Despite our characterizing the particle in terms of winding numbers, there is nothing in this argument that necessitates η being an integer.

Compare our upgraded 2D exchange to the old-fashioned definition. If we carry out the exchange $(x_1, x_2) \rightarrow (x_2, x_1)$ [shown in Fig. 3(c)] then the formal definition of exchange tells us that the wavefunction should be identical (for bosons) or pick up a minus sign (fermions). However, the new version merely tells us that $\phi = \pi$, resulting in a phase factor of $\Phi(\phi) = e^{i\eta\pi}$. The two versions of exchange are only identical for the special cases that (a) we have $\eta =$ even integer, when we recover the expected exchange behaviour for bosons or (b) we have $\eta =$ odd integer, when we recover fermion exchange. However, this analysis shows that there are many more possible values of η in 2D, since it doesn't have to be an integer. We are therefore not tied simply to bosons and fermions, the freedom to choose η means we can have any exchange statistics, and particles with such statistics are called (Abelian) anyons.⁷

5. Gauge theory

QSLs are often named with a mathematical group (e.g. a Z_2 spin liquid, or U(1) spin liquid). This is related to the properties of the QSL's gauge structure, considered in this section through a number of examples.

Example 1: Consider complex scalar field theory, which is a model built from continuous complex-number valued fields $\psi(x)$ that are described by a Hamiltonian

$$\hat{H} = (\partial_x \psi)^{\dagger} (\partial_x \psi) + m^2 \psi^{\dagger} \psi, \qquad (3)$$

where $\partial_x \equiv \partial/\partial x$ and *m* is a constant. This theory has an internal symmetry known as a global phase symmetry, or global U(1) symmetry. This is a shorthand for the observation that with the replacement $\psi(x) \to \psi(x)e^{i\alpha}$, the Hamiltonian *H* does not change.⁸ This is a global transformation in that the phase changes by the same amount (α) at every point.

What if we attempt to change the phase *locally*, i.e. differently at each point? That is, change the phase by an amount $\alpha(x)$ that depends in some arbitrary manner on position. Under the local transformation, the derivatives change as follows:

$$\partial_x \psi(x) \to \partial_x \left[\psi(x) \mathrm{e}^{\mathrm{i}\alpha(x)} \right] = \mathrm{e}^{\mathrm{i}\alpha(x)} \partial_x \psi(x) + \psi(x) \mathrm{e}^{\mathrm{i}\alpha(x)} \mathrm{i}\partial_x \alpha(x)$$
$$= \mathrm{e}^{\mathrm{i}\alpha(x)} \left\{ \partial_x + \mathrm{i} \left[\partial_x \alpha(x) \right] \right\} \psi(x). \tag{4}$$

Similarly $\partial_x \psi^{\dagger}(x) \to e^{-i\alpha(x)} \{\partial_x - i [\partial_x \alpha(x)]\} \psi^{\dagger}(x)$, and the first term in the Hamiltonian becomes

$$(\partial_x \psi)^{\dagger} (\partial_x \psi) - \mathrm{i} \psi^{\dagger} (\partial_x \alpha) (\partial_x \psi) + \mathrm{i} \psi (\partial_x \psi^{\dagger}) (\partial_x \alpha) + \psi^{\dagger} \psi (\partial_x \alpha) (\partial_x \alpha), \tag{5}$$

which is not what we started with. Perhaps unsurprisingly, the theory is not therefore invariant with respect to local phase transformations. However, to fix things we introduce a new gauge field whose job is to cancel out the effect of the change in internal variable $\alpha(x)$ with position. This enters into the Hamiltonian as a covariant field derivative D_x , defined by $D_x\psi(x) = \partial_x\psi(x) - iqA^x(x)\psi(x)$, where q is a coupling constant. We define the gauge field such that if we change the phase by angle $\alpha(x)$, the components of the gauge field transform according to $A^x(x) \to A^x(x) + \frac{1}{q}\partial_x\alpha(x)$. Now if $\psi(x) \to \psi(x)e^{i\alpha(x)}$, then

$$D_x\psi = (\partial_x - iqA^x)\psi \to e^{i\alpha}(\partial_x\psi) + i\psi e^{i\alpha}(\partial_x\alpha) - iqA^x\psi e^{i\alpha} - i\psi e^{i\alpha}(\partial_x\alpha) = e^{i\alpha}D_x\psi.$$
(6)

This property makes the whole Hamiltonian invariant under local phase changes if we

⁷Having only two spatial dimensions is vital to the argument. In 3D all type-I processes are topologically identical since they are all deformable into paths where the particles don't move. Similarly all type-II processes are topologically identical and may be reduced to a simple exchange of particles. This reduction occurs because the extra dimension allows us to move the paths past each other in the third dimension, shrinking all loops to zero.

⁸This transformation generates the elements of the Lie group U(1), hence the name.

replace ordinary derivatives by covariant ones: $\hat{H} = (D^x \psi)^{\dagger} (D_x \psi) + m^2 \psi^{\dagger} \psi$, since now with $D_x \psi \to e^{i\alpha} D_x \psi$, the first term is invariant.

The message of these manipulations is that if the phase $\alpha(x)$ is a function of x then, in order to guarantee local phase invariance, a new gauge field $A^x(x)$ is required to form the covariant field derivative. Furthermore, we recognise this field as akin to the electromagnetic gauge field, where one of its features is that, without changing E and B, the gauge field can be changed by an arbitrary amount $A^x(x) \to A^x(x) + \partial_x \chi(x)$. If we identify $\chi(x)$ with $\alpha(x)/q$, then the transformation demanded by our argument is simply electromagnetic gauge invariance, confirming that A^x has the usual properties of a gauge field as defined in electromagnetism. We say that the gauge field has its own dynamics, which in electromagnetism are governed by the Maxwell equations, written in terms of the (3+1)-dimensional field A^{μ} .

Example 2: In Section 7 we will introduce a model of spinon particles on a lattice represented in terms of the operator \hat{f}_i^{\dagger} , which creates a spinon at site *i*. This is similar to the continuous model in Example 1, except that position is now a discrete variable. We again demand invariance under a local U(1) phase transformation $\hat{f}_i \rightarrow \hat{f}_i e^{i\phi_i}$, which changes the phase at each lattice point *i*. To ensure this we again introduce a gauge field, this time $\bar{\chi}_{ij}e^{i\theta_{ij}}$, where θ_{ij} is a phase that depends on the bond between sites *i* and *j*. The spinons and the gauge field interact according to a Hamiltonian of the form $\hat{H} = \hat{f}_i^{\dagger} \hat{f}_j (\bar{\chi}_{ji}e^{i\theta_{ij}})$, so if the gauge field transforms according to $\theta_{ij} \rightarrow \theta_{ij} + \phi_i - \phi_j$, then the Hamiltonian is symmetric under local U(1) transformations of the spinon operators. We will see later that this gives us a U(1) gauge theory of spinons.

Example 3: Z_2 (more often written \mathbb{Z}_2 in the mathematics literature) is a group with two elements: 1 and -1. For the Kitaev model on a lattice (Section 8), we identify an operator \hat{c}_i and demand invariance under a local transformation $\hat{c}_i \to W_i \hat{c}_i$, where W_i is an arbitrary function that outputs 1 or -1 at the *i*th site. We identify a gauge field \hat{u}_{ij} that depends on two positions *i* and *j*, and a Hamiltonian of the form $\hat{H} = (\hat{u}_{jk})\hat{c}_j\hat{c}_k$. If \hat{u}_{ij} transforms according to $\hat{u}_{ij} \to W_i \hat{u}_{ij} W_j^{-1}$, then this allows the theory to be locally Z_2 invariant.

Example 4: Our discussion in the next section relies on another Z_2 gauge field S_{ij} , where positions are specified by two lattice sites *i* and *j*. Here we will again transform via the local transformation $S_{ij} \to W_i S_{ij} W_j^{-1}$ in the same way as Example 3. This transformation leaves the Hamiltonian invariant, as we shall now discuss.

6. Our second spin liquid: the toric code

There is a class of models that allows us to get a handle on some of the ingredients we've met so far: we start with the **toric code** and then (Section 6.2) consider a simplified version of this model. We shall see how these models give rise to spin-liquid ground states with topological ground-state degeneracy, anyonic excitations and a gauge-field structure based on the group Z_2 .

The toric code [12,21] describes spins on a 2D lattice with periodic boundary conditions. The boundary conditions are a key feature: by identifying each of the two perpendicular directions on the lattice, we obtain a space with the topology of a torus (hence the name of the model). An example torus is shown in Fig. 4(a). This surface has two holes: one through the empty middle of the doughnut, and one through the tubular part (filled with dough in the edible version). The important thing to note is that there are three main sorts of paths on the torus [see Fig. 4(a)]: A-type and



Figure 4. (a) A torus with three different sorts of paths on its surface. [Figure reproduced from Ref. [4], reprinted with permission of Oxford University Press.] (b) Square lattice in which spins lie in the spaces between lattice points. Example spin positions are shown in green in the upper left. An example plaquette p and star s are also shown. (c) Flipping a connected string of spins aligned along $\pm x$ (red dots) excites two e anyon (or star) excitations at the ends of the string (labelled s_1 and s_2). (d) Flipping spins aligned along $\pm z$ along a string perpendicular to the bonds (blue dots) creates two m anyon (or plaquette) excitations (labelled p_1 and p_2). (e) Top: the action of the star operator on a path through the lattice flips spins (purple), giving a loop of flipped spins on the *dual* lattice (green). Bottom: a line of flipped spins (blue) makes an A-type path on the torus. (f) Configuration of bonds and plaquettes that results from applying periodic boundary conditions to a lattice of 4 points. Arabic numerals label bonds, roman numerals label plaquettes.

B-type paths wrap the two different holes that characterise the torus, while C-type paths wrap neither of the holes. It is impossible to smoothly deform one type of path into a different one (that is, they are topologically distinct: we would need to cut the paths and glue them back together to do this).

The lattice in the toric code model is decorated by s = 1/2 spins which are, rather unusually, situated on the centres of the *bonds*, rather than on the vertices of the lattice, as one would usually expect, [the spins are represented by green dots in Fig. 4(b)]. We shall mostly work in a basis of spins directed up and down along z, and operate on a spin between sites i and j with Pauli sigma operators using the usual rules $\hat{\sigma}_{ij}^{z}|\uparrow_{ij}\rangle = |\uparrow_{ij}\rangle, \hat{\sigma}_{ij}^{z}|\downarrow_{ij}\rangle = -|\downarrow_{ij}\rangle, \hat{\sigma}_{ij}^{x}|\uparrow_{ij}\rangle = |\downarrow_{ij}\rangle$ and $\hat{\sigma}_{ij}^{x}|\downarrow_{ij}\rangle = |\uparrow_{ij}\rangle$. There are two contributions to the Hamiltonian [see Fig. 4(b)]: (i) from each square

There are two contributions to the Hamiltonian [see Fig. 4(b)]: (i) from each square plaquette (labelled with an index p) on the lattice, and (ii) from each star, which is the cross-shaped structure formed from the nearest-neighbour bonds of lattice site s. The contribution from a plaquette is given by $\hat{P}_p = \prod_{ij \in p} \hat{\sigma}_{ij}^z$, which is to say a product of the z-components of spins on the bonds around a plaquette. This takes values ± 1 . The contribution from a star is given by $\hat{R}_s = \prod_{ij \in s} \hat{\sigma}_{ij}^x$, (i.e. operate with the x-sigma matrix, flipping spins in our σ^z basis). This contribution also takes values ± 1 (most easily seen by transforming to a basis of spins directed along $\pm x$). Each of these two operators commutes amongst themselves and, as can be checked, they commute with each other. We put the two contributions together to form the Hamiltonian for the toric code

$$\hat{H} = -g\sum_{p}\hat{P}_{p} - t\sum_{s}\hat{R}_{s},\tag{7}$$

where g and t are positive constants.

6.1. Ground state and excitations

On one level, the toric code is quite a simple model whose properties can be guessed. We can immediately spot a candidate ground state $|0\rangle$: if we can find a wavefunction such that all eigenvalues $R_s = P_p = 1$ we obtain the lowest-possible energy. Consider plaquette q, to get $P_q = 1$ we must have an even number of down spins. (There are, of course, several ways to arrange this.) We can also see the role of entanglement. Consider the ground state expressed in a basis of eigenstates of the $\hat{\sigma}^x$ operator. Each spin is shared between two stars. To make sure a star has $R_s = 1$ there must be an even number of down (along-x) spins in the star. Moreover, in order that all stars have $R_s = 1$, these down spins must be must be arranged in closed loops. (Try it and see!) The ground state does not have a preferred spin direction and so is magnetically disordered and is an example of a quantum spin liquid.

Excitations from the ground state have some plaquettes or stars contributing -1 to the Hamiltonian. Excitations with $P_p = -1$ are known as vortices. These are bosons, they cost energy 2g to produce, and are also known as magnetic (m) particles (or sometimes as visons). The other sort of excitation is a single negative star $R_s = -1$, which is also a boson, this time costing energy 2t, and is known as an electric (e)charge. Neither m nor e can be created locally as individual excitations using a single operator, as we do with particles (where $c_i^{\dagger}|0\rangle = |1_i\rangle$). For the e particles, for example, the best we can do is flip a string of spins, which creates *two* localised e particles: one at each lattice site at the end of the string of flipped spins [Fig. 4(c)]. This is because each spin is shared between two lattice sites, and hence shared by two stars. The *m* particles are also created in pairs, occurring at plaquette centres at the ends of strings of flipped bonds, this time as shown in Fig. 4(d). Again this is due to each spin being shared between two plaquettes. This non-locality of excitations leads to anyon statistics, although in slightly more subtle form to what we've seen so far. Although *e* and *m* are both bosonic, moving an *e* around an *m* gives a π phase shift, resulting in a negative sign. We say that the particles have *mutual statistics*, specifically we call them mutual *semions* in this context. Additionally, the composite particle ε made up of a boson *e* and a boson *m* is actually a fermion.

Using these ideas, we can deduce some basic properties of the toric-code ground state more formally. Working in a basis of eigenstates of σ^z operators, we specify spin configurations by writing a wavefunction $|\{S_{ij}\}\rangle$, where $\{S_{ij}\}$ are a list of the eigenstates of each operator σ_{ij}^z for the bond between sites *i* and *j*. The ground state $|0\rangle$ will be built from a superposition of these spin states $|\{S_{ij}\}\rangle$. To find $|0\rangle$, we require for the plaquette term that $\hat{P}_p|0\rangle = |0\rangle$. This is only possible for a state of the form

$$|0\rangle = \sum c_s |\{S_{ij}\}\rangle, \qquad (8)$$

where the c_s s are constants and the sum is constrained to be over those spin configurations that feature no vortices (i.e. where $P_p = 1$).

In this σ^z basis, the star operators act on an $|\{S_{ij}\}\rangle$ to flip spins. If we start with a state with no vortices and act with the star operator on sites that lie on paths around the lattice then we create a pattern of down spins on an up-spin background. (The star operator never creates vortices owing to the commutation of the two operations.) If we visualise this by linking the down spins via the *dual lattice* (i.e. the lattice formed from points in the middle of each lattice point on the original lattice), the result is always closed loops of down spins, of the sort shown in green in the example in Fig. 4(e). A ground state is formed from the superposition of all such vortex-free states comprising closed loops of down spins on the dual lattice. These all contribute with equal coefficients $c_s = c$ in eqn 8. The resulting state gives $\hat{R}_s |0\rangle = |0\rangle$, confirming that we have a ground state.

Although this accounts for one ground state, there are actually three more degenerate states, making four in total. To access the others, let's consider the Wilson loop operator $\hat{U}(\mathcal{C})$, which multiplies the spins around a closed contour \mathcal{C} . The operator is given explicitly by

$$\hat{U}(\mathcal{C}) = \hat{S}_{ij}\hat{S}_{jk}...\hat{S}_{li},\tag{9}$$

where the indices are selected to take us around the contour. The eigenvalue $U(\mathcal{C})$, which can take values ± 1 , is called the Z_2 flux through \mathcal{C} . [For example, the flux from a vortex excitation (i.e. $P_p = -1$ on a plaquette p) is U = -1.] Since the star operator flips *pairs* of spins along any contour through the lattice, it must commute with $\hat{U}(\mathcal{C})$. Now, if we take the path around the two distinct paths that wrap a torus $[\ell_1 = A \text{ or}$ $\ell_2 = B$ in Fig. 4(a)], then we must have $U(\ell_i) = \pm 1$. [The ground-state configuration we formed in the last paragraph has U(A) = U(B) = 1.]

The four degenerate ground states come from the possibility of adding loops of down spins that wrap all the way around the torus in either of those two distinct paths that link the torus [i.e. paths A or B in Fig. 4(c)]. An example of an A-type loop of down spins is shown in Fig. 4(e). A B-type contour that wraps the torus must cross this line of down spins and yield U(B) = -1 [we also have U(A) = +1]. Note that there is no way to create the down-spin loop from the application of the star operators, and also that the state containing just this loop has no vortices. Compared to the state with no down-spin loops that wrap the torus, this state will give distinct patterns of down spins on the dual lattice once we start working on it with the star operator. As a result, the state formed by all down-spin loops in an U(A) = 1, U(B) = -1, system yields a second degenerate ground state. The other degenerate states have U(A) = -1, U(B) = +1[i.e. a vertical line of flipped spins on Fig. 4(e)] and U(A) = U(B) = -1 [i.e. both a vertical and horizontal line of flipped spins].

The resulting 4-fold degeneracy is a *topological* degeneracy, meaning that it occurs for all lattices on the torus. In the language of flux, we might call U(A) the flux through one hole of the torus and U(B) the flux through the other. States with U = 1 have no flux through the torus while states with U = -1 have one unit of flux threading it.⁹ The degeneracy is very robust, as we can see if we treat the star term as a small perturbation to the vortex term by assuming $g \gg t$ in eqn 7. Recall from degenerate perturbation theory that to compute the splitting caused by a perturbation, we need to identify a series of operators that takes the ground state up to an excited state (or states) and then returns it back to another state in the ground-state manifold. The only way this can occur here [1] using the star operator is for it to create excited stated by flipping spins all of the way around the torus. For an $L \times L$ lattice, this requires L excitations. Perturbation theory tells us that any energy gap caused by such a perturbation is of order $\Delta E \approx t^L/g^{L-1}$, but as we tend to the thermodynamic limit of a system $L \to \infty$, we find $\Delta E \to 0$. We conclude that the topological degeneracy is robust. In fact, the degeneracies in our toric code are exact, and not split at all by t^{L}/q^{L-1} terms, although such a splitting would apply in a more generic model.

6.2. Counting states in Z_2 theory

This subsection, and the following one, are a tutorial given in terms of a simplified model theory to explicitly compute the degeneracy via a demonstration of the gauge structure. You can skip to Section 7 at this stage if desired.

Let's simplify the toric code model further and consider only the plaquette term [1] with spins acted on by operators $\hat{S}_{ij} = \sigma_{ij}^{z}$ (that is, we set t = 0 in eqn 7). The Hamiltonian is now simply $\hat{H} = -g \sum_{p} \hat{P}_{p}$, where $\hat{P}_{p} = \prod_{ij \in p} \hat{S}_{ij}$ now operates on the spins in a plaquette p only. For definiteness take the number of sites on the lattice to be n = 4 so we have only a single plaquette p = I. In the absence of boundary conditions, the spins on this plaquette can adopt $2^{n} = 16$ configurations, half giving $P_{\rm I} = -1$.

Now apply the all-important boundary conditions and get Fig 4(f) with 2n = 8 spins on the resulting bonds (labelled 1-8 in the figure), and therefore $2^{2n} = 256$ possible spin configurations. However, the description in the Hamiltonian $\hat{H} = -g \sum_p \hat{P}_p$ is not given at the level of spins, but at the higher level of plaquettes. By seeking to describe the system's states at this higher level, we'll see that a description in terms of P_p can be visualised as a pattern of fluxes that thread the lattice. Compared to the model without boundary conditions, the periodic boundaries provide three additional plaquettes, making four in total [labelled I, II, III and IV in Fig. 4(f)]. However, these four plaquettes are not independent, since each spin is shared between two plaquettes.

⁹In fact, for a genus g Riemann surface, we can put zero or one unit of flux through each of the 2g holes of the surface, leading to a 2^{2g} degenerate ground state.

Plaquette	Spins in	No. of free choices to
p	the plaquette	obtain a given P_p
Ι	$1\ 2\ 3\ 4$	8
II	$1 \ 3 \ 5 \ 8$	2
III	$2\ 4\ 6\ 7$	2
IV	$5\ 6\ 7\ 8$	0

Table 1. Decorate the periodic lattice in Fig. 4 with spins to realise a pattern of $(P_{\rm I}, P_{\rm II}, P_{\rm III}, P_{\rm IV})$, in the order $p = \rm I$, II, III, IV. There are 8 choices to obtain $P_{\rm I}$, but because spins are shared, these are constrained for the other plaquettes. The table shows that there are $8 \times 2 \times 2 = 32$ spin configurations that give each pattern of fluxes $\{P_p\}$. (We check that 32×8 different $\{P_p\}$ gives our 256 possible spin configurations.)

If we multiply all of the plaquette operators together, each spin features twice in the product, and since $P_p = \pm 1$, we have the constraint $\prod_p P_p = 1$. A consequence is that there are only 3 possible values of energy $E = \sum_{p=1}^{IV} P_p$, namely 4, 0 and -4. There is one set of $\{P_p\} = (P_1, P_{II}, P_{III}, P_{IV})$ that leads to E = 4, one set that gives E = -4, and six sets with E = 0, making 8 different allowed sets (called "patterns of flux" below) in total. This is a lot less than the 256 possible spins states! In fact, consulting Table 1, we find that there are $8 \times 2 \times 2 = 32$ ways of arranging spins to get each allowed choice of $\{P_p\}$.

Takeaway: there are $2^{2n} = 256$ configurations of spins in the model. There are only 8 unique sets of $\{P_p\}$. That is, there are 8 distinct patterns of flux through the lattice. It's therefore impossible to uniquely label spin configurations with a list of P_p .

Although there are 8 patterns of flux, we don't actually know how many states exist in the Hilbert space of the system when it is described in terms of fluxes. This is because both degeneracy and gauge equivalence will cause there to be several spin states that give an equivalent pattern of fluxes.

Evaluating the balance between degeneracy and gauge equivalence is the subject of the next subsection. However, if formal gauge theories are not to your taste, here is the direct argument: our simplified model has t = 0 in eqn 7. We might ask how the energies of these sets of states are split for $0 < t/g \ll 1$. The answer is that all the exact eigenstates for t, g > 0 can be labelled by the positions of the e and m particles, plus the values of the fluxes around the two loops that encircle the torus. Here we have 8 configurations of m particles (i.e. one with no plaquettes occupied, one with all four plaquettes occupied, and six with two of the plaquettes occupied), and for each set of m particles there are 8 possible configurations of e particles (i.e. one with no sites occupied, one with all four sites occupied, and six with two of the sites occupied). For each of these 64 configuration of m and e particles we've seen that there are 4 possible configurations of flux that wrap the torus. So the total number of states is $8 \times 8 \times 4 = 256$ as expected. For t, g > 0 these states must exist as 4-fold degenerate levels, since the values of the flux don't influence the energy of the states.

6.3. Introducing gauges

Many of the 256 spin states are gauge equivalent. Recall that "gauge equivalence" is a statement that several different configurations specified by a gauge field can correspond to the same state in the Hilbert space of a system. We see this if we make a local transformation: the transformed system gives the same Hamiltonian and observables. For our problem, the Hamiltonian is given in terms of the configuration of plaquettes, which therefore determines the Hilbert space. We will call these states in the Hilbert space, which can be visualised in terms of the patterns of fluxes through the plaquettes, the *physical states* for this problem. Since many of the underlying spin states correspond to the same pattern of fluxes, we can therefore treat the specific description of spins as a gauge field. To summarise, on putting the spins on bonds and specifying the configurations using the the plaquette operators \hat{P}_p in our Hamiltonian, rather than the spins, we obtain a redundancy of description, and hence we say that a gauge structure *emerges*.¹⁰

In our model, S_{ij} is the gauge field and we'll see that several states described by different S_{ij} correspond the same state in the Hilbert space of the model when described in terms of plaquette operators. Technically, we can group spin states into collections called *gauge-equivalent classes*, where all of the states in such a class are related by a local transformation and so represent the same physical state. [1] As a result, states in a Hilbert space are in a 1-1 relation with the number of gaugeequivalent classes.

The gauge group for our model is Z_2 , which is to say specifically that a local transformation W_i of the spin field S_{ij} leaves the Hamiltonian invariant. Here the local transformation W_i is an arbitrary function that takes values ± 1 at the different lattice sites. Two spin fields S_{ij} and \tilde{S}_{ij} are gauge equivalent if they are related by a local transformation W_i via $\tilde{S}_{ij} = W_i S_{ij} W_j^{-1}$. You can see an example of how this works by starting with a state \bar{S}_{ij} , selecting some site k and setting $W_k = -1$ and all of the other W_i to be +1. This flips spins in a star shape around site k to give a configuration \tilde{S}_{ij} . Since two of the flipped bonds feature in each plaquette surrounding k, the Hamiltonian $\hat{H} = -g \sum_p \hat{P}_p$ is gauge invariant: you get the same value of energy, no matter whether you use \bar{S}_{ij} or \tilde{S}_{ij} . Moreover, you can make $2^n = 16$ different W_i functions on our lattice, so this is the number of possible local transformations.

In order to work out if two states on a lattice are really the same physical state (i.e. if they are gauge equivalent), an important role is played by the Wilson-loop operator $\hat{U}(\mathcal{C})$ from eqn 9. The key is that $U(\mathcal{C})$ itself is gauge invariant, so states that give different values of $U(\mathcal{C})$ cannot be related by an arbitrary local transformation. Therefore if we know the pattern of fluxes provided by evaluating $U(\mathcal{C})$ around each plaquette, then we have the physical description of a state. Note that if \mathcal{C} is a single plaquette p then $\hat{U}(\mathcal{C}) \equiv \hat{P}_p$, and so the sets $\{P_p\}$ pick up the flux through each plaquette and label states in a way that's invariant with respect to local transformations. In short: a different $\{P_p\}$ means a different physical state. However, the same $\{P_p\}$ doesn't guarantee that all of the states are gauge equivalent (i.e. the same physical state) because we don't yet know how many states are distinct states that are degenerate. Returning to our example, we have that a number (32 in our case) of spins states will give the same pattern of fluxes described by a given $\{P_p\}$.

To finally compute the degeneracy, we need a technical argument [1], which goes as follows: There are $2^n = 16$ different possible local transformations, but there are two special transformations that don't change the spin configuration, namely $W_i = 1$ and $W_i = -1$, for all i. A (non-trivial) consequence of the existence of these two special transformations is that the 16 gauge transformations can actually only create $2^n/2 =$

¹⁰The situation here is subtly different to that of electromagnetism, where there is no direct experimental access to the gauge field. Here we could imagine making measurements of the spins in a real system via magnetic susceptibility. However, by treating the physics at the (Hamiltonian) level of plaquettes and, equivalently, fluxes, there's a sense in which we're treating the spins as a microscopic description that we can't access, and working instead with what emerges at the more coarse-grained level of fluxes.

8 gauge-equivalent configurations, where the 2 in the denominator is the number of these special transformations. With this key fact in hand, we can say that there are $2^{2n}/(2^n/2) = 2 \times 2^n = 32$ gauge-equivalent classes [i.e. total number of spin states (256), divided by number of gauge-equivalent configurations (8)] and therefore 32 different physical states in the Hilbert space. So out of 256 total spin configurations, only 32 are distinct physical states. To label the flux states we use the Z_2 flux through a plaquette provided by the eigenstates of the \hat{P}_p operator. We saw that there are only $2^n/2 = 8$ different values of $\{P_p\}$ and so only 8 distinct patterns of flux through the system. However, if each state is 4-fold degenerate we are saved, since then the different flux patterns give the $8 \times 4 = 32$ physical states in the Hilbert space.

Takeaway: There are 32 physical states in the Hilbert space of the model. Each of these states corresponds to 8 gauge-equivalent states (accounting for 256 spin configurations). There are 8 different patterns of flux encoded by $\{P_p\}$. Each of the 8 flux patterns corresponds to 4 physically-distinct degenerate states.¹¹

Let's pause and consider what we've learnt from this technical discussion. The simplified toric-code model has a spin-liquid ground state which respects rotational spin symmetry and is highly entangled. The model is invariant with respect to local Z_2 transformations, giving it a rich gauge structure. The topology of the lattice, via the gauge structure, causes all states to be 4-fold degenerate. Finally, we might congratulate ourselves by getting this far with a quotation from X-G. Wen [1]:

If you feel the definition of Z_2 gauge theory is formal and the resulting theory is strange, then you get the point.

7. Spinons

An important excitation found in most spin-liquid models is the spinon. The spinon is an s = 1/2, neutral fermion and therefore (compared, e.g., to the s = 1/2, charged electron) an example of a fractional excitation. The idea of fracturing of the quantum numbers of an electron is easiest to see in one dimension [18], which demonstrates the remarkable feature that the apparently fundamental properties of the electron: spin and charge, can break into two. We imagine a set of electrons arranged along a onedimensional line as shown in Fig. 5(a), with electron spins aligned up-down-up-down, so the overall spin of the system is zero. If we remove an electron from this system then we leave behind a hole, which can be thought of as an excited state of charge (known as a holon in this context). If we move an electron along the line into the empty space without changing its spin, then we see that the hole is mobile. This has a consequence: it leaves two like spins as neighbours forming a s = 1/2 spinon excitation. As we slide electrons around we see that the holon can move independently of the spinon. Similarly by flipping a pair of spins we can move the spinon around. The spin excitation and the hole excitation are independent: we have spin-charge separation.

In the purely magnetic case (i.e. without the charge excitation), we attempt to measure the dispersion by flipping a spin [Fig. 5(b)] to create a s = 1 magnon excitation with wavevector q. This is not stable on the 1D chain and so falls apart into a pair of s = 1/2 spinons that propagate via pairs of spin flips with wavevectors q_1 and q_2 , where $q = q_1 + q_2$. We therefore expect to measure (in an inelastic neutron scattering measurement) a continuum of possible excitations [the grey shading in Fig. 5(c)] lying

¹¹Since each of the physically distinct states can be expressed as 8 gauge equivalent configurations of spins, there are 32 spin configurations giving each pattern of fluxes, as we found above.



Figure 5. (a) Spin and charge separate in a one-dimensional electron system. A spin up electron is removed, leaving a hole behind. As the hole moves down the chain it leaves behind it a spin excitation (circled). (b) In a spin chain we create a s = 1 magnon by flipping a spin. This is not a stable excitation but can split into a pair of s = 1/2 spinons (circled) that can propagate via flips of spin pairs. (c) Spinons pairs give rise to a continuum of excitations in the chain (grey area) lying between the magnon dispersion (upper curve) and the single spinon dispersion (lower curve).

between the dispersion curves of a single spinon and the single magnon. It is the characteristic continuum of excitations that tells us we have fractional excitations.

7.1. Representing spins as fermions

It is rather surprising that we can form spinons from the spins in a magnet. In this section we discuss how this is possible within a mean-field picture of magnetism.¹² The idea will be to fractionalise the spins into parts or "partons". If we guess how to make the split correctly, the partons might represent the true excitations of the system. As has been commented [1,12], some practitioners are uncomfortable with this approach, but we shall willingly suspend disbelief for the moment.

Let's return to two dimensions and to the Heisenberg antiferromagnet with Hamiltonian $\hat{H} = -\sum_{\langle ij \rangle} J_{ij} \hat{S}_i \cdot \hat{S}_j$ and make a *mean-field approximation* [4]. This essentially causes a spin to sit in the constant average magnetic field of its ordered neighbours, allowing us to compute its dynamics. The usual recipe to do this replaces the original Hamiltonian by

$$\hat{H}_{\rm mf} = -\sum_{\langle ij\rangle} J_{ij} \left[\langle \hat{\boldsymbol{S}}_i \rangle \cdot \hat{\boldsymbol{S}}_j + \hat{\boldsymbol{S}}_i \cdot \langle \hat{\boldsymbol{S}}_j \rangle - \langle \hat{\boldsymbol{S}}_i \rangle \cdot \langle \hat{\boldsymbol{S}}_j \rangle \right].$$
(10)

where $\langle \hat{\boldsymbol{S}}_i \rangle = \langle \Phi_{\rm mf} | \hat{\boldsymbol{S}}_i | \Phi_{\rm mf} \rangle$ and $| \Phi_{\rm mf} \rangle$ is the mean-field ground state of the system. Although this approximation represents the spirit of this section, this prescription won't actually work for a QSL, since the states with which we're interested have $\langle \boldsymbol{S}_i \rangle = 0$, as we expect them to be disordered.

To make progress, we perform a transformation trick that allows us to represent the spin as a chargeless fermion.¹³ We introduce spinon operators $\hat{f}_{i\alpha}$ where *i* is a

 $^{^{12}}$ Although our discussion is given in terms of fractionalising spins to form fermionic spinons, it is also possible to choose to form bosonic spinons via the use of the Schwinger boson representation of spin [19]. Such an approach has the appealing feature that some states can be described in terms of the Bose-Einstein condensation of such spinons [22].

¹³It's worth noting that such transformations have a rich history in many-body physics. Jordan and Wigner spotted that a single spin state can be thought of as an empty, or singly-occupied fermion state using a mapping $|\uparrow\rangle \equiv f^{\dagger}|0\rangle$ and $|\downarrow\rangle \equiv |0\rangle$. However, to deal with more than one spin they must add a phase factor, called a string, to each fermion. [19]

site on the lattice and $\alpha = 1, 2$ [1,14]. These are two-component, fermionic s = 1/2, charge-neutral operators (sometimes called Abrikosov fermions). There are two sorts of these (f_{i1} spinons and f_{i2} spinons). A spin operator is represented by spinons as

$$\hat{\boldsymbol{S}}_{i} = \frac{1}{2} \hat{f}_{i\alpha}^{\dagger} \boldsymbol{\sigma}_{\alpha\beta} \hat{f}_{i\beta} = \frac{1}{2} \begin{pmatrix} \hat{f}_{i1}^{\dagger} & \hat{f}_{i2}^{\dagger} \end{pmatrix} \begin{pmatrix} \boldsymbol{\sigma}_{11} & \boldsymbol{\sigma}_{12} \\ \boldsymbol{\sigma}_{21} & \boldsymbol{\sigma}_{22} \end{pmatrix} \begin{pmatrix} \hat{f}_{i1} \\ \hat{f}_{i2} \end{pmatrix}.$$
 (11)

In these equations, sums over the Greek indices, which take the values 1 and 2, are implied.

The result of the spinon transformation is that the magnetic Hamiltonian becomes $\hat{H} = \frac{1}{2} \sum_{\langle ij \rangle} J_{ij} \hat{f}^{\dagger}_{i\alpha} \hat{f}_{j\alpha} \hat{f}^{\dagger}_{j\beta} \hat{f}_{i\beta} + \text{const.}$ To recap, we have recast the magnetic Hamiltonian in terms of chargeless spinons which yields an expression in the spirit of the tight-binding model of eqn 1, that can be interpreted in terms of fermions hopping between sites. By making this transformation we've allowed there to be four possible spinon states per site $(|f_{i1}f_{i2}\rangle = |00\rangle, |01\rangle, |10\rangle, |11\rangle)$, where previously there were only two spin states ($|\uparrow\rangle$ and $|\downarrow\rangle$). To make sure there's one fermion per site we should, strictly speaking, set the constraint $\hat{f}^{\dagger}_{i1}\hat{f}_{i1} + \hat{f}^{\dagger}_{i2}\hat{f}_{i2} = 1$, which reduces the possible states down to $|10\rangle \equiv |\uparrow\rangle$ and $|01\rangle \equiv |\downarrow\rangle$ only.

We can now make progress by imposing a mean-field approximation on the spinon model [1,14]. Recall that this involves taking an average of combinations of operators. This is carried out here by relaxing the constraint on the number of fermions per site, such that we set $\langle \hat{f}_{i\alpha}^{\dagger} \hat{f}_{i\alpha} \rangle = 1$, which says that the *average* number of fermions per site is now one. In practice this is done using a Lagrange multiplier V_i , leading to a new term $V_i(f_{i\alpha}^{\dagger} f_{i\alpha} - 1)$ in the Hamiltonian. We enact the mean-field approximation, which now amounts to taking averages of combinations of pairs of spinon operators, and throws up the answer

$$\hat{H}_{\rm mf} = \frac{1}{2} \sum_{\langle ij \rangle} J_{ij} \left[\hat{f}^{\dagger}_{i\alpha} \hat{f}_{j\alpha} \chi_{ji} + h.c. - |\chi_{ij}|^2 \right] + \sum_i V_i (\hat{f}^{\dagger}_{i\alpha} \hat{f}_{i\alpha} - 1), \tag{12}$$

where *h.c.* is the Hermitian conjugate of the preceding term. The new quantity $\chi_{ij} = \langle \hat{f}_{i\alpha}^{\dagger} \hat{f}_{j\alpha} \rangle = \langle \hat{f}_{i1}^{\dagger} \hat{f}_{j1} \rangle + \langle \hat{f}_{i2}^{\dagger} \hat{f}_{j2} \rangle$ is key here. It measures the amplitude for either sort of fermion to hop from site *j* to site *i*. The quantities χ_{ij} and V_i don't change if we rotate spins, so we can conclude that the ground state also has this rotational invariance. The model can now yield spin-liquid ground states.

From the previous discussion, we would conclude that there is a disordered, correlated ground state whose excitations are spinons. However, to get a reliable picture that maps to the Hilbert space of the original spin model, we must also include the possibility of low-energy fluctuations in χ_{ij} by saying $\chi_{ij} = \bar{\chi}_{ij} e^{-i\theta_{ij}}$ and obtain

$$\hat{H}_{\rm mf} = \frac{1}{2} \sum_{\langle ij \rangle} J_{ij} \left[\hat{f}_{i\alpha}^{\dagger} \hat{f}_{j\alpha} \bar{\chi}_{ji} \mathrm{e}^{-\mathrm{i}\theta_{ij}} + h.c. \right] + \sum_{i} V_i (\hat{f}_{i\alpha}^{\dagger} \hat{f}_{i\alpha} - 1), \tag{13}$$

We see that if we make the local phase transformation $\theta_{ij} \rightarrow \theta_{ij} + \phi_i - \phi_j$ and $\tilde{f}_i \rightarrow \tilde{f}_i e^{i\phi_i}$ the Hamiltonian is unchanged. We have discovered the U(1) gauge structure of the model. We identify the V_i and θ_{ij} , which interact with the spinons in the same form, as components of a gauge field, and conclude that the excitations of the theory are spinons *coupled* to this U(1) gauge field. To make further progress we need to specify the form of χ_{ij} , which then allows us to derive the dispersion of the excitations for specific QSL models. This can result in a range of states including RVB-type models and a several Z_2 spin liquids [1].

For now, we can recap the main features of this approach to finding QSLs: in order to use a mean-field theory to describe magnetically-disordered states we have needed to fractionalise the spins before taking the averages on which the mean-field technique relies. The result is a theory that predicts magnetically disordered ground states with neutral s = 1/2 fermion excitations. On allowing fluctuations in the averaged quantities we obtain a gauge structure, where we interpret these fluctuations quantum mechanically as bosonic excitations in the gauge field. In short, we split up the spins at the start, and we use bosonic glue (via the gauge-field component V_i) to stick them back together! This results in a Hamiltonian model that can support a spin liquid ground state, with predictions of its excitations.

This general approach can be extended to more complicated gauge structures to make a range of spin liquids states. We expect stable mean-field QSLs to have a non-zero energy gap against gauge field fluctuations. The stable QSLs always contain neutral s = 1/2 spinons with short range interactions between them. The mean-field approach predicts four families of spin liquid in (2+1) dimensions, characterised by the nature of the energy gap between the ground state and the spinon and gauge excitations. (See Ref. [1] for the full story.)

Rigid spin liquids are topologically ordered, which means both spinon excitations and the excitations in the gauge field have non-zero energy gaps. The gaps lead the QSLs to be stable since there are no low-energy excitations to destabilise the states. Examples include Z_2 -gapped liquids and chiral liquids.

Bose spin liquids are a class characterised by spinon excitations with an energy gap, and gapless U(1) gauge-boson excitations. These are not stable in (2+1) dimensions. Fermi spin liquids have gapless s = 1/2 spinon excitations with short-range interactions between them. The name derives from Fermi liquids, which are characterised by their gapless electronic excitations. Examples include Z_2 -linear-, Z_2 -quadratic-, Z_2 -gapless liquids, where names refer to features of the spinon dispersion.

Algebraic spin liquids have gapless excitations, but these are neither free bosons nor free fermions. Excitations are massless fermions coupled to U(1) gauge field. Examples include the U(1)-linear liquid.

If you've succeeded in suspending disbelief until now, disbelief is surely flooding back! Should we trust any of these manipulations? Arguably, the jury is still out, but the approach does suggest the features of some of the possible states that might be out there for discovery. Ultimately experiment will be the arbiter and we shall discuss the status of experiments at the end of the article. For now we shall turn to a different means of realising a QSL that, unlike the mean-field picture, does not rely on approximations.

8. The Kitaev model

For many of the formative years of the subject, the search for QSLs was inspired by the successes such as the FQH fluid and BCS superconductivity, that suggested that identifying wavefunctions, rather than Hamiltonians, was the way to discover new states of matter. Although this approach has indeed been transformative, there was also a great deal of value in the later identification of Hamiltonian models of QSLs. A major step forward has been the formulation of the Kitaev model [21]. This is not least



Figure 6. The Kitaev model. (a) Honeycomb lattice with directional bonds. (b) To solve the model, we define a set of particle operators attached to each vertex of the lattice. (c) A plaquette in the lattice leads to the definition of the operator \hat{W}_p . (d) Dispersion for $J_x = J_y = J_z$ showing Dirac cones. (e) The triangular phase diagram given in terms of (J_x, J_y, J_z) . (f) For $J_z \gg J_x, J_y$, we create superspins along the z bonds and recreate the toric code lattice. A sample plaquette and star is shown.

because it is solvable, based on a simple magnetic Hamiltonian that has a spin-liquid ground state. The solution, which involves the use of operators representing Majorana fermions (a hypothetical fermion that is its own antiparticle), looks potentially rather unfamiliar, but the result is less so: it leads to a Z_2 gauge field and vortex excitations, just like we saw in Section 6. Even better, in one limit of interactions the system gives us back the toric code model exactly.

The Kitaev model is based on a planar honeycomb lattice with spins on each vertex [Fig. 6(a)]. The links are labelled x, y, z and the Hamiltonian is written

$$\hat{H} = -J_x \sum_x \hat{\sigma}_i^x \hat{\sigma}_j^x - J_y \sum_y \hat{\sigma}_i^y \hat{\sigma}_j^y - J_z \sum_z \hat{\sigma}_i^z \hat{\sigma}_j^z.$$
(14)

To solve the model we introduce an operator that operates on hexagonal plaquette p shown in Fig. 6(c) via $\hat{W}_p = \hat{\sigma}_1^x \hat{\sigma}_2^y \hat{\sigma}_3^z \hat{\sigma}_4^x \hat{\sigma}_5^y \hat{\sigma}_6^z$, with eigenvalues $W_p = \pm 1$. The \hat{W}_p for each p commutes so we can describe the Hilbert space in terms of a set $\{W_p\}$. It will turn out that the ground states has $W_p = 1$ for all plaquettes.

Although this is progress, the full solution requires swapping spins for a new set of operators as we did in the last section. Fascinatingly, these are the operators describing Majorana fermions. Conventional fermions can be created (with operator \hat{d}^{\dagger}) an annihilated (\hat{d}) , with the constraints $\hat{d}^{\dagger}\hat{d}^{\dagger} = 0$ and $\hat{d}\hat{d} = 0$, and anticommutator $\hat{d}^{\dagger}\hat{d} + \hat{d}\hat{d}^{\dagger} = 0$. A Majorana fermion is its own antiparticle, which means that conjugating the charge of a Majorana has no effect. In terms of operators this means we represent Majoranas as real-valued. Therefore, if we have two species of Majorana (*a* and b particles) then we could split up a fermion as $\hat{d} = (\hat{a} + i\hat{b})/2$, where the Majorana operators for this representation are \hat{a} and \hat{b} . It follows that the Majoranas have the properties $\hat{a} = \hat{a}^{\dagger}$ and $\hat{b} = \hat{b}^{\dagger}$ (i.e. the creation and annihilation operators are the same), $\hat{a}^2 = \hat{b}^2 = 1$ (i.e. applying the operator twice has no effect) and $\hat{a}\hat{b} + \hat{b}\hat{a} = 0$ (the operators anticommute with each other).

8.1. The Kitaev model solution and its gauge structure

For our problem [21], we will need to introduce four Majorana particles at each vertex of the lattice [Fig. 6(b)]. At each vertex we therefore define a set of Majorana operators $\hat{b}^x, \hat{b}^y, \hat{b}^z, \hat{c}$, and we write spin operators in terms of these as $\hat{\sigma}^x = i\hat{b}^x\hat{c}, \hat{\sigma}^y = i\hat{b}^y\hat{c}$, and $\hat{\sigma}^z = i\hat{b}^z\hat{c}$. The new operators square to unity and they anticommute in the same way that the σ matrices do. However, we find that we also need an extra constraint: identify operator $\hat{D} = \hat{b}^x \hat{b}^y \hat{b}^z \hat{c}$ (which commutes with the spin operators) and, with the stipulation that $D = \pm 1$, we guarantee a consistent theory.¹⁴ So why do we need four Majoranas? We'll give the *c*-field a local Z_2 invariance, and the *b*-fields will be the gauge fields that guarantee this. This will give us a QSL ground state with fermion excitations that derives from the *c*-field, along with vortex excitations that will arise from the flux of the gauge field.

With the proposed set of transformations, the Kitaev Hamiltonian becomes

$$\hat{H} = \frac{i}{4} \sum_{\langle jk \rangle} 2J_{\alpha} \left(i\hat{b}_{j}^{\alpha} \hat{b}_{k}^{\alpha} \right) \hat{c}_{j} \hat{c}_{k}.$$
(15)

The first thing to notice is that this looks like a tight-binding model for the *c*-particles. Moreover, the shape of this Hamiltonian has the *c*-type Majoranas interacting with an effective field $\hat{b}\hat{b}$. In this case, the combination $\hat{b}_j^{\alpha}\hat{b}_k^{\alpha} = \hat{u}_{jk}$, with α determined by the direction of the bond [i.e. $\alpha = \alpha(i, j)$], is a Z_2 gauge field, since $u_{jk} = \pm 1.^{15}$ The operator \hat{u}_{jk} commutes with the Hamiltonian, so we can split up the solutions according to the values of u_{jk} using the product around a plaquette $\hat{W}_p = \prod \hat{u}_{jk}$, which we now see evaluates the flux of the gauge field through a plaquette.

The ground state has $W_p = 1$ for all plaquettes. This state is a magneticallydisordered (i.e. spin-rotation invariant), spin-liquid phase. The Hamiltonian can then be diagonalised to reveal the dispersion for the *c*-field excitations of $E(\mathbf{k}) = \pm 2 |J_x e^{i\mathbf{k}\cdot\mathbf{n}_1} + J_y e^{i\mathbf{k}\cdot\mathbf{n}_2} + J_z|$, where $\mathbf{n}_{1,2} = (\pm\sqrt{3}/2, 3/2)$. As usual, we characterise the material in terms of the gap against excitations. An example of the spectrum for $J_x = J_y = J_z$ is shown in Fig. 6(d) and has several sets of excitations existing from E = 0, each with a striking linear dispersion. The points in *k*-space from which they emerge are known as Dirac points, where the dispersion resembles that of highly relativistic particles (i.e. with linear dispersion E = |p|c, where *c* is the speed of light).

We can always find some fermion solution with E = 0 as long as the exchange constants obey triangle inequalities: $|J_x| \leq |J_y| + |J_z|$ or $|J_y| \leq |J_z| + |J_x|$ or $|J_z| \leq |J_x| + |J_y|$. When these aren't obeyed, the spectrum has a gap. We can draw a triangular phase diagram of gapped and gapless phases as shown in Fig. 6(e). The phases A_i are

 $^{^{14}}$ The point here is that the Hilbert space dimension for a single spin is 2 but for four majoranas it is 4, so we need to project out two spurious states per site.

¹⁵To see this, note that the local transformation in this case would send $c_j \to W_j c_j$ and then a change in the gauge field of $u_{jk} \to W_j u_{jk} W_k^{-1}$ cancels out this effect.

completely gapped, while the phase B is gapless. In addition to the fermion excitations, we have $W_p = -1$ excitations in the gauge field. These are bosonic vortices.

Finally, as promised, in the limit of strong J_z -coupling we recover the toric code Hamiltonian. Physically, this is because the strong z-links tie the spins together to form effective spins $|1\rangle = |\uparrow\uparrow\rangle$ and $|2\rangle = |\downarrow\downarrow\rangle$ which decorate the edges of a square lattice. This is shown in Fig. 6(f), where a sample plaquette and star are indicated.

9. Realising a QSL in materials

We've seen several ways of describing different putative QSLs, but the question remains whether the state is realised in real materials. As we said at the start of this article, there is a wide-ranging literature claiming QSL behaviour in a large number of materials. However, it's probably fair to say that we do not have a case where a consensus has been reached in favour of the occurrence of the QSL state. Despite this we have reason to be hopeful. We have dealt with (2+1) dimensions in our discussion so far, but the analogous one-dimensional (1D) case of the spin-Luttinger liquid is agreed to be realised in some spin-chain materials, in which direct measurements of spinons have been made, most directly using inelastic neutron scattering (INS). The spin-Luttinger regime is characterised by a lack of long-range magnetic order (LRO), algebraically-decaying spin correlations and, most vividly, a continuum of spinon excitations [Fig. 5(c)]. Here it is unwise to say a material *is* a spin-Luttinger liquid, rather that it behaves as a spin-Luttinger liquid under certain conditions. The 1D chain will be characterised by an intrachain interaction J. If the chains have some small interchain interaction J' then the material might well order at a very low temperature $T_{\rm N}$ determined by J'. In order to observe 1D physics we seek to be in a regime that $T_{\rm N} \ll T \ll J$. This means we stay away from magnetic order, but the temperature is low enough that the collective behaviour is promoted by J. (For $T \gg J$ excitations will start to resemble single uncorrelated spin flips.)

A QSL will be realised in more than one dimension. Although we have dealt with 2D in our discussion, there are also proposals for 3D QSL states. To realise a QSL we seek to suppress magnetic order in materials with strong spin-spin interactions. Low dimensionality is good for this, although many 2D materials enjoy enough coupling in the third dimension to promote order at low temperature. Magnetic order can also be suppressed by frustrating magnetic interactions. (When this occurs, small interactions can tip the system over into LRO at a low-enough T, so these must be borne in mind.) We also want to promote quantum fluctuations and for this s = 1/2 spins are most effective. Other possibilities involve promoting higher-order exchange interactions, such as next-nearest neighbour couplings.

The experimental identification of a QSL is a knotty problem, as discussed in Refs [23,24]. Clearly we seek a system without magnetic order down to low temperature. A lack of order can be difficult to establish, but sensitive local probes such as NMR and muon-spin spectroscopy (μ^+ SR) are often used to search for weak signs of order, or spin dynamics characteristic of disordered spin-liquid states. We have stressed that the excitation spectrum is key to the characterisation of a QSL system. Perhaps the most important question is whether there is a gap against excitations and there are several ways of investigating this that are sensitive to the density of states, including thermodynamic, scattering and spectroscopic probes. Beyond this, a determination of the excitation spectrum might be expected to supply the smoking-gun evidence. INS remains a key method here, but requires samples that are relatively large com-



Figure 7. (a) Structure of Herbertsmithite showing 2D Kagome layers [28]. (Copyright Springer, Nature.) (b) κ -(ET)₂Cu₂(CN)₃ is based on a 2D triangular lattice of BEDT-TTF dimers. (c) The 3D pyrochlore structure [29] (Copyright American Physical Society.) (d) α -RuCl₃ showing the hexagonal lattice realised by the Ru³⁺ ions [30]. (Copyright AAAS.) Reprinted with permission.

pared to those required for other measurements. There is also the problem that some nuclei have large neutron-capture cross sections, or the material contains hydrogen or other nuclei that cause a large degree of incoherent scattering. Nevertheless, we look for fractionalised quasiparticles which, like spinons in 1D, give broad responses in scattering as they correspond to the creation of several particles. In the absence of momentum-resolved scattering measurements, there are predictions for the temperature and magnetic field dependence of properties such as heat capacity, spin relaxation or transport. In the case of the latter, although QSLs should strictly be insulators, there are possibilities for the state to couple to charges, allowing an additional handle to probe the system.

We have stressed that the key to the QSL problem is the special pattern of quantum entanglements that underlie the physics. Finding a means of evaluating these could prove the key to this subject. The entanglement entropy S of the many-particle ground state is useful here. For a gapped system it follows the form $S = cL - \Gamma + ...$, where L in the first term is the size of the system, while the second term quantifies long-range entanglements [23]. This latter part is only non-zero in a topological phase (for a Z_2 QSL, $\Gamma = \ln 2$). The bad news, however, is that there isn't yet a known way of measuring this quantity.

Below we briefly introduce four different candidate systems that have been much discussed as possible approximate realisations of a QSL state. This represents tip of the iceberg in terms of the extensive literature, which now includes lots of distinct avenues: for example possible QSLs in cold-atom experiments [25]. Further discussion of our four chosen candidates, including the many relevant references, can be found in the detailed reviews of Refs. [12–14,26,27].

Herbertsmithite or $ZnCu_3(OH)_6Cl_2$ is a mineral, originally discovered in a

Chilean mine, but since synthesised in the laboratory. It comprises s = 1/2 spins on a frustrated 2D kagome lattice interacting with an exchange of J = 180 K [Fig. 7(a)]. The kagome lattice is formed from corner sharing triangles (as distinct from edge-sharing triangles of the triangular lattice), resulting in a larger degeneracy and potential for realising QSL states. No evidence of long-range magnetic order has been found in Herbertsmithite down to 50 mK. Specific heat and some NMR measurements suggest gapless behaviour, although a spin gap has also been identified on the basis of Knightshift measurements. Large crystals have enabled neutron scattering, but the broad continuum of magnetic excitations found was not one predicted by spin liquid theory. In fact, the low-energy excitation spectrum appears dominated by impurity spins that originate from sites between the kagome planes. Although these sites should host nonmagnetic Zn^{2+} , they appear to contain enough magnetic Cu^{2+} to form a significant population of orphaned impurity moments (estimated to be 5-10%). This raises the question of how robust a QSL might be with respect to disorder, which is a question that is pertinent to all of the cases discussed here. The presence of perturbing anisotropic (Dzyaloshinskii-Moriya) interactions between intrinsic kagome spins has also been found using electron spin resonance, which could have a large effect on the ground state [31].

 κ -(ET)₂Cu₂(CN)₃ is formed from a layered structure built from structural dimers of planar BEDT-TTF molecules, sandwiched between layers of Cu₂(CN)₃. Each dimer gives rise to a magnetic s = 1/2 spin that form a 2D triangular lattice with a large exchange J > 200 K [Fig. 7(b)]. This material is part of a family of materials that show superconductivity and magnetic order, along with a number of spin-liquid candidates. For κ -(ET)₂Cu₂(CN)₃ there is no evidence of magnetic order from μ +SR or NMR. The combination of constant low-temperature magnetic susceptibility χ , linear specific heat (i.e. like a metal $C = \gamma T$) and a resulting Fermi-gas like Wilson ratio $W (\propto \chi T/C)$, often used to characterise interactions in electron systems) was suggestive of a gapless phase with spin-carrying excitations. However conductivity and μ +SR both suggested the presence of an energy gap. Owing to the tunability of molecular materials, several other systems based on similar building blocks also show some promise.

 $R_2M_2O_7$ (R = rare earth, M = transition metal such as Ti or Sn) are a well know frustrated series of 3D materials based on spins in a 3D arrangement of corner-sharing tetrahedra, called the pyrochlore lattice [Fig. 7(c)], which is known to exhibit a high degree of frustration. Strong spin-orbit interactions and crystal field splitting can result in systems described by an effective s = 1/2 theory. Some of these materials are well-described by classical theories (e.g. Ho₂Ti₂O₇ and Dy₂Ti₂O₇). For example, in $Dy_2Ti_2O_7$ the crystal-field anisotropy constrains the magnetic moments to lie along the local $\langle 111 \rangle$ axes (i.e. directly in our out of the tetrahedron) and there is an effective local ferromagnetic coupling between these moments, as well as long-range dipolar couplings, which turn out to be important. As a result of this combination of interactions and the local anisotropy, below about 1 K the system settles into a disordered *spin-ice state*, which is often described as a classical spin liquid (i.e. a highly correlated magnetic system that avoids magnetic order). Spin ice is characterised by a '2-in 2-out' spin configuration (meaning that two spins point in and two spins point out of each tetrahedron), analogous to proton displacement vectors in Pauling's model of hydrogen disorder in water ice. The excitations in spin ice are created by reversing a single spin, which produces a pair of effective magnetic monopoles which can move independently through the lattice and interact with an emergent gauge field, but remain connected by a topological Dirac string of flipped spins between them. In contrast, strong quantum effects have been found in $R_2 Ti_2 O_7$ with R = Yb, Er and Tb and in $Pr_2Zr_2O_7$. Of these Yb is the most studied, having been suggested to be a quantum version of spin ice which features frustrated Ising interaction. Such a model is thought to support a U(1) QSL in some limits of its parameters. This picture has some experimental support (e.g. it features a "pinch point" feature in its neutron spectrum, predicted for classical spin ice owing to dipolar spin correlations). However, there is also good evidence for a phase transition to a ferromagnetic ground state in high-quality samples, and the unusual neutron spectrum has been suggested to arise from an interplay of types of conventional order [32]. This is another case where there is a degree of sample dependence, suggesting a role for disorder.

Finally, there are several candidate materials that might be described by the **Kitaev** Hamiltonian. Although the Kitaev Hamiltonian, with bond-dependent Ising interactions, appears somewhat contrived, it was suggested that partially-filled t_{2q} levels in an octahedral environment with strong spin-orbit coupling could realise the model. The idea was originally to look at materials containing Ir^{4+} , which has an effective j = 1/2moment, owing to the orbital degeneracy being lifted in favour of a Kramers doublet. Studies of honeycomb α -Na₂IrO₃ and α -Li₂IrO₃ revealed magnetic order, but also indicated that the bond-directed Kitaev interactions were, to some extent, realised. Iridium is well known for its high neutron capture cross section, making INS difficult and so the most-studied Kitaev candidate is not an Ir-hosting material, but α -RuCl₃ [Fig. 7(d)]. Here j = 1/2 ruthenium ions form a honeycomb lattice between Cl planes. This system also orders $(T_{\rm N} = 15 \text{ K})$, and seems to host conventional Heisenberg interactions in addition to bond-directed ones. However, there is evidence from Raman scattering for a fermionic continuum and magnetic order seems to be fully suppressed by a strong applied field, suggesting that the interactions might be tuned to QSL behaviour. As a result, the system is sometimes called a proximate spin liquid. So is the honeycomb lattice really sufficient for the Kitaev model? It is possible that the spin liquid regime might only be stable in a tiny region and the interpretation of data on these materials has been hotly debated. Owing to the large amount of work on this subject at the time of writing, this remains an area to watch.

None of these systems uncontroversially show QSL behaviour, but all have features suggesting they might be close. As with the 1D materials the key to observation of this behavior will be looking at the realm of applicability of the models to the materials. If the QSL states are realised in some form, and it seems likely that they are, evidence will likely continue to accumulate slowly as the vast parameter space of possible materials is surveyed. One essential ingredient to consider is disorder. This is an inevitable feature of condensed matter systems and will be of relevance here. It is perhaps worth keeping in mind its dual role in the related field of quantum Hall physics: the FQH fluid was only observed in relatively clean samples, but disorder is also believed to be necessary in order to observe the Hall plateaux. It is possible disorder might adopt a similar dual role in the story of the QSL, or at least resolve some of the contradictory experimental evidence that is found in the current literature.

10. Conclusion

The study of quantum magnetism started with doubts about the existence of the antiferromagnet but, through careful experiment and thoughtful theoretical description, the field now boasts a wealth of established results that demonstrate the fractionalisation of excitations in 1D and the existence of topological excitations related to vortices.



Advances in materials preparation, including those towards the ability to engineer magnets from tuneable building blocks, also provide hope that hitherto purely-theoretical models might be realised in material systems. The quest to realise a quantum spin liquid remains a much sought-after goal that goes well beyond simply finding a material that fails to order at low temperatures. The aim of this article has been to give a sense of what is at stake here beyond a lack of magnetic order: namely a macroscopic manifestation of a very specific type of quantum entanglement that presents us with a zoology of elementary excitations over and above those that fit comfortably within Landau's "Standard Model" of condensed matter physics. From this point of view, the investigation of QSLs could help us in the next phase of determining the ordering principles that quantum mechanics imposes on our material world.

Acknowledgements

I am grateful to the following colleagues for their many helpful comments on and corrections to this article: Stephen Blundell, Nathan Bentley, Theo Breeze, Claudio Castelnovo, Matjaž Gomilšek, Thomas Hicken, Ben Huddart, Matt Jones, Samuel Ladd and Francis Pratt [who also supplied Fig. 7(b)]. Particular thanks are due to John Chalker who made detailed comments on two drafts of the review.

Biography

Tom Lancaster is Professor of Physics at Durham University. His research focuses on the physics of magnetism in one and two dimensions along, more recently, with the physics of topological objects in magnetism. He is the coauthor, with Stephen Blundell, of Quantum Field Theory for the Gifted Amateur.

References

- [1] Wen, X-G. Quantum Field Theory of Many-body Systems, OUP, Oxford (2004).
- [2] Landau, LD. and Lifshitz, EM. Statistical Physics, Pergamon, Oxford (1980).
- [3] Anderson, PW. Basic Notions of Condensed Matter Physics, Westview, Boulder (1984).
- [4] Lancaster, T. and Blundell, SJ. Quantum Field Theory for the Gifted Amateur, Oxford University Press, Oxford (2014).
- [5] Singleton, J. Band Theory and Electronic Properties of Solids, OUP, Oxford (2001).
- [6] Lieb E., Schultz T., and Mattis D. Ann. Phys. 16, 407 (1961).
- [7] Blundell, SJ. Magnetism in Condensed Matter, Oxford University Press, Oxford (2001).
- [8] Feynman, RP., Leighton RB., and Sands M. The Feynman Lectures on Physics, Vol. I, Pearson Addison Wesley, San Francisco (2006).
- [9] Anderson PW. Concepts in Solids, World Scientific, London (1997).
- [10] Lhuillier C. Frustrated Quantum Magnets arXiv:cond-mat/0502464 (2002).
- [11] Anderson PW. Mater. Res. Bull. 8, 153 (1973).
- [12] Savary L. and Balents L. Rep. Prog. Phys. 80, 016502 (2017).
- [13] Broholm C., Cava RJ., Kivelson SA., Nocera DG., Norman MR. and Senthil T. Science 367, eaay0668 (2020).
- [14] Zhou Y., Kanoda K. and Ng T-K. Rev. Mod. Phys. 89, 025003 (2017).
- [15] Manousakis E. Rev. Mod. Phys. 63, 1 (1991).
- [16] Boothroyd A.T. Contemp. Phys. 63, 305 (2022).
- [17] Lancaster T. Contemp Phys. 60, 246 (2019).
- [18] Giamarchi T. Quantum Physics in One Dimension, OUP, Oxford (2003).
- [19] Coleman P. Introduction to Many body Physics CUP, Cambridge (2015).
- [20] Aitchison IJR. and Mavromatos NE. Contemp. Phys. 32, 219 (1991).
- [21] Kitaev A. Annals of Physics **321** 2, (2006); Kitaev A. and Laumann C. arXiv:0904.2771v1 (2009).
- [22] Xu C. and Sachdev S. Phys. Rev. B 79, 064405 (2009).
- [23] Knolle J. and Moessner R. Annu. Rev. Condens. Matter Phys. 10 451 (2019).
- [24] Khatua J., Sana B., Zorko A., Gomilsek M., Sethupathi K., Ramachandra Rao MS., Baenitz M., Schmidt B. and Khuntia P. arXiv:2310.15071 (2023).
- [25] Semeghini G., Levine A., Keesling H., Ebadi S., Wang TT., Bluvstein D., Verresen R., Pichler H., Kalinowski M., Samajdar R., Omran A., Sachdev S., Vishwanath A., Greiner M., Vuletić V. and Lukin MD. Science **374**, 1247 (2021).
- [26] Clark L. and Abdeldaim AH. Annu. Rev. Mater. Res. 51, 495 (2021).
- [27] Wen J., Yu S-L., Li S., Yu W. and Li J-X. npj Quantum Materials 4, 12 (2019).
- [28] Khuntia P., Velazquez M., Barthélemy Q., Bert F., Kermarrec E., Legros A., Bernu B., Messio L., Zorko A. and Mendels P. Nature Phys. 16, 469 (2020).
- [29] Elhajal M., Canals B., Sunyer R. and Lacroix C. Phys. Rev. B 71, 094420 (2005).
- [30] Yokoi T., Ma S., Kasahara Y., Kasahara S., Shibauchi T., Kurita N., Tanaka H., Nasu J., Motome Y., Hickey C., Trebst S. and Matsuda Y. Science 373, 568 (2021).
- [31] Zorko A., Nellutla S., van Tol J., Brunel LC., Bert F., Duc F., Trombe J-C., de Vries MA., Harrison A. and Mendels P. Phys. Rev. Lett. 101, 026405 (2008).
- [32] Scheie A., Kindervater J., Zhang S., Changlani HJ., Sala G., Ehlers G., Heinemann A., Tucker GS., Koohpayeh SM. and Broholm C. PNAS 177, 27245 (2020).



Citation on deposit: Lancaster, T. (2023). Quantum spin liquids. Contemporary Physics, <u>https://doi.org/10.1080/00107514.2023.2284522</u>

For final citation and metadata, visit Durham Research Online URL: <u>https://durham-</u> repository.worktribe.com/output/2047134

Copyright statement: This accepted manuscript is licensed under the Creative Commons Attribution 4.0 licence. https://creativecommons.org/licenses/by/4.0/