Mocking quantum mechanics: Semiclassical and machine-learning approaches to frustrated magnetism

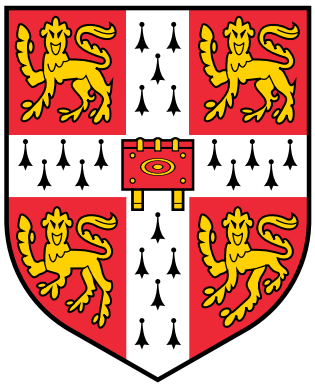

Attila Szabó

Trinity College
University of Cambridge

Supervisor: Prof. Claudio Castelnovo

This dissertation is submitted for the degree of Doctor of Philosophy

October 2020

Declaration

I hereby declare that except where specific reference is made to the work of others, the contents of this dissertation are original and have not been submitted in whole or in part for consideration for any other degree or qualification in this, or any other university. This dissertation is my own work and contains nothing which is the outcome of work done in collaboration with others, except as specified in the text and Acknowledgements. This dissertation contains fewer than 60 000 words including appendices, bibliography, footnotes, tables and equations and has fewer than 150 figures.

Several chapters of this dissertation are based on research published in, or submitted to, scientific journals:

Chapter 2 A. Szabó, C. Castelnovo, *Seeing beyond the light: Vison and photon electrodynamics in quantum spin ice.* Phys. Rev. B **100,** 014417 (2019, Editors' Suggestion)

Chapter 3 A. Szabó, G. Goldstein, C. Castelnovo, A. M. Tsvelik, *Generalised route to effective field theories for quantum systems with local constraints.* Phys. Rev. B **100**, 085113 (2019, Editors' Suggestion)

Chapter 4 A. Szabó, C. Castelnovo, *Neural network wave functions and the sign problem.* Phys. Rev. Research **2**, 033075 (2020)

Chapter 5 M. J. Pearce, K. Götze, A. Szabó, T. S. Sikkenk, M. R. Lees, A. T. Boothroyd, D. Prabhakaran, C. Castelnovo, P. A. Goddard, *Monopole density and antiferromagnetic domain control in spin-ice iridates.* arXiv:2102.04483 (2021, under review in Nature Communications)

Chapter 6 A. Szabó, U. Schneider, *Non-power-law universality in one-dimensional quasicrystals.* Phys. Rev. B **98**, 134201 (2018) and
A. Szabó, U. Schneider, *Mixed spectra and partially extended states in a two-dimensional quasiperiodic model.* Phys. Rev. B **101**, 014205 (2020)

Attila Szabó
October 2020

Acknowledgements

Gratitude is not only the greatest of virtues, but the parent of all others.

— Marcus Tullius Cicero

First and foremost, I thank my supervisor, *Claudio Castelnovo.* He has been the greatest PhD supervisor and mentor I could wish for. Beyond everything I learnt from him about physics (which is a lot!), his enthusiasm for science, positive attitude towards research, conviviality, and kindness has always been a driving force on my research journey thus far. I have learnt from him the value of exploring interesting questions for their own sake, not worrying too much whether a project "pays off," and the ways of finding true enjoyment in science; but however many lessons I'll learn from him, I think there will always be more reason to look up to him as a scientist, teacher, and good person.

I also owe gratitude to all other collaborators, without whom the research in this thesis wouldn't have been possible. Namely, I thank *Ulrich Schneider* for introducing me to the exotic world of quasicrystals and localisation, suggesting the challenges whose solutions are discussed in chapter 6, and continuing to work with me after my master's project to bring those solutions to completion. I acknowledge *Garry Goldstein* and *Alexei Tsvelik* for suggesting the soft-constraint ideas underlying the work in chapter 3, and the research groups of *D. Prabhakaran* and *Paul Goddard,* and especially *Matt Pearce* and *Kathrin Götze,* for the excellent experimental work that motivated the studies in chapter 5 and is discussed at length in the same. I am also indebted to the exceptional intellectual environment of the TCM group in the Cavendish; conversations by The Coffee Machine have shaped my view of physics and the world more than I would recognise, and supplied me with more ideas than I could remember or properly acknowledge. Special thanks goes to *Max* and *Andrea* for their careful proofreading of the thesis and their thoughtful comments on grammar, style, and science.

Hálás vagyok mindazoknak, akik mostanáig kísértek utamon és felkeltették bennem a fizika iránti érdeklődést: Elsősorban szüleimnek és tágabb családomnak, akik örömmel vállaltak minden nehézséget, amivel egy tehetséges, külföldre szakadó gyerek felnevelése jár, akik mindig ott voltak, amikor szükségem volt rájuk, és mérhetetlen szeretettel vittek előre az úton. Gimnáziumi fizikatanáromnak, Simon Péter*nek, akinek elhivatottsága és*

fizikaszeretete mindenkit magával ragad és a legjobbat hozza ki belőlük; sosem felejtem el, hogy már jóval ezelőtt lehetőséget adott arra, hogy vele együtt írjak egy könyvet! A pécsi és budapesti olimpiai szakkörök vezetőinek, Kotek László*nak,* Vigh Máté*nak,* Vankó Péter*nek és* Honyek Gyula*nak, akik kielégítették tudásvágyamat és a lehető legjobb alapokat adták leendő pályámhoz.*

I shall thank all those who have supported me on the way to this point and those who kindled and nurtured my interest in physics: First of all, my parents and broader family, who happily put up with all difficulties of having a gifted child and later an émigré, for always being there when I needed them and guiding me along the way with immense love and care. My high-school physics teacher, *Péter Simon,* whose dedication and love of his subject infects everyone around him and brings the best out of them; indeed, he gave me an occasion to write a book with him well before this one! The leaders of advanced study groups, *László Kotek, Máté Vigh, Péter Vankó,* and *Gyula Honyek,* who satisfied my thirst for knowledge and gave the best possible preparation for my eventual career. The fellows and graduate students of Trinity College, Cambridge, and especially *Dima Khmelnitskii,* whose supervisions were truly without comparison.

Man shall not live by science alone, so I wish to thank all those who made the past seven years in Cambridge a time to always remember fondly. The students and postdocs of TCM, friends rather than coworkers: *Petr, Ollie, Max, Alice, Beñat, Jan, Stephen, Eze, Adam, Chris, Andrea, Angela, Ivona, Benji, Danny, Elis, Tobias, James, Ariel, Jonathan,…* Friends from Trinity: *Matthew, Filip, Katya, Will, D'Arcy, Georg, Maria, Alex, Herschel, Irene, Leonardo, Jaan, Petr, Nat, Steph, Neil,…* The Cambridge Hungarians, a piece of home in a strange land: *Peti, Márk, Menyus, Balu, Sztanka, Rozi, Gerry, Bence* (all of you!), *Robi, Rebeka, Olivér, Áron, Marcell, Andris, KDP, Botond, Price, Roland, Bálint, Dávid, Tomi,…* All those I have forgotten to list – the names should go on for pages! I thank you all for being by my side, showing all the affection one could wish for, and providing just the right amount of distraction to balance work and life; I am happy for all pub nights, parties, long walks along the Cam, and all-night chats in kitchens with you. I am indebted to the *University Counselling Service, Juliet,* and *Gillian* for guiding me to the bright side of life the times when everything seemed dark, and helping me discover how to stay there. Last, but definitely not least, I thank everyone at *St Clement's, Cambridge,* this beautifully quirky and convivial community joined by the love of God and old-fashioned English, for being a true spiritual home during my time in Cambridge.

Soli Deo gloria.

Abstract

Many-body quantum mechanics is the fundamental theory behind many areas of modern science, such as condensed-matter physics, nuclear physics, and quantum chemistry. It is also notoriously hard: the classical picture of particles with well-defined positions and velocities is replaced by an intricate interference pattern between all their possible trajectories, captured by the quantum wave function. The exponentially large information content of wave functions makes direct simulation of large, strongly interacting quantum systems impossible, and necessitates strategies to manage the complexity in an analytically or computationally tractable manner.

The bulk of this thesis explores two such strategies in the context of quantum spin liquids. In these materials, competition between incompatible interactions results in robust, massive entanglement, down to zero temperature. Such ground states give rise to a range of exotic behaviour, such as topological order and fractionalised excitations: understanding these is a central challenge in the physics of strongly correlated materials.

Several quantum-spin-liquid phases are underpinned by strict local conservation laws, which give rise to lattice gauge theories with exotic quasiparticle excitations, such as emergent photons or magnetic monopoles. We developed a systematic approach [1], based on a large-S bosonisation formalism, to extract gauge-theoretic descriptions from such constrained Hamiltonians automatically, and thus make them amenable to the powerful techniques of quantum field theory. The same field theories also allow us to simulate quantum-spin-liquid systems semiclassically, that is, to replace spin-$1/2$ operators with classical "compass needles," removing the computational complexity of quantum entanglement without losing the physical behaviour. We demonstrated this approach on quantum spin ice, a paradigmatic and experimentally relevant model of quantum spin liquids and, by simulating it on unprecedented large system sizes, obtained novel insights about its quasiparticles [2].

The success of neural networks in a range of machine-learning problems makes them a natural candidate for representing highly entangled quantum states, allowing in principle an accurate simulation of large, challenging quantum systems with modest computational resources. However, most current approaches using such neural quantum states suffer from the infamous Monte-Carlo sign problem, making deep neural networks unable to learn ground states in antiferromagnetic and fermionic systems. We studied the

possible origins of this sign problem and proposed a neural-network ansatz that is able to overcome the sign problem for unfrustrated antiferromagnets [3].

In addition to this main theme, I have been part of an experiment–theory collaboration on understanding the unique magnetoresistance properties of the classical frustrated magnet $Ho_2Ir_2O_7$ [4]. We discovered a mechanism by which antiferromagnetic domains can be coupled to an external magnetic field via an intercalated spin-ice system: such control is highly desirable for spintronics applications. We have also identified scattering channels through which magnetic monopoles give rise to a significant contribution to the resistivity of $Ho_2Ir_2O_7$: this allows us to directly measure their density in a simple and flexible experiment.

Finally, I report studies on the localisation properties of quasicrystals; namely, the discovery of thermodynamic universality not described by the usual power laws in one-dimensional quasiperiodic models [5], and of a two-dimensional quasicrystal in which localised and partially extended states coexist without a mobility edge [6].

Contents

1

Quantum spin liquids: an introduction

Quantum spin liquids (QSL) are exotic phases of matter, arising in a range of frustrated magnetic systems, in which quantum zero-point fluctuations prevent the formation of long-range magnetic order down to zero temperature. This "quantum disorder" is associated with massive many-body entanglement in the ground state, which brings about unusual physical properties, such as non-local excitations, fractionalised quasiparticles, and topological order [7]. In this chapter, I discuss two paradigms of describing QSL phases: the first is based on perturbative processes between the degenerate ground states of frustrated Ising models; the second exploits the relationship between Heisenberg and Hubbard models to describe the former in terms of fractionalised spin-1/2 fermions. Both approaches allow for a description of the frustrated spin model in terms of lattice gauge theories, which manifestly encode topological order and fractionalisation. Finally, I review the chemistry of rare-earth pyrochlore materials, an important experimental platform for realising spin liquids, as well as the classical spin-ice model that describes the pyrochlores $Dy_2Ti_2O_7$ and $Ho_2Ti_2O_7$ and serves as the background of chapters 2 and 5.

1.1 Classical and perturbative spin liquids

The loftier the building, the deeper must the foundations be laid.

— Thomas à Kempis

Perhaps the oldest and simplest example of a frustrated magnet is the nearest-neighbour antiferromagnetic Ising model on the triangular lattice: it is geometrically frustrated because the three spins around each triangle cannot all be made antiparallel. In the thermodynamic limit, this implies that a third of all interactions cannot be satisfied, even in the ground state, where each triangle has precisely one frustrated link. It is intuitively clear that there is a huge number of ways to arrange these frustrated links: for

instance, a finite fraction of spins in a typical ground-state configuration would have three neighbours of either sign, allowing them to be flipped at no energy cost. These and other, larger-scale, rearrangements of spins are allowed even at zero temperature; as a result, the model never develops conventional magnetic order. This is associated with an exponential number of degenerate[†] ground states, which leads to an extensive zero-temperature entropy,[‡] in apparent violation of the third law of thermodynamics. Nevertheless, the situation is very different from a collection of non-interacting spins exhibiting uncorrelated disorder: interactions restrict ground-state configurations through an extensive number of local constraints, resulting in nontrivial correlations (which are often long-range despite the originally short-range interactions), as well as topological order that is only observable through measurements that span the entire system. These exotic properties mark a new phase of matter, qualitatively different from both ordered magnets and trivial paramagnets, called *Ising* or *classical spin liquids,* in analogy with liquids that are neither long-range ordered (as solids are) nor consist of independent particles (as gases do).

The extensive ground-state entropy of spin liquids has a profound effect on its excitations. Let us focus on spin liquids where the ground state is characterised by an extensive number of local constraints such that flipping any individual spin – the usual elementary excitation of an Ising model – violates more than one constraint.[§] These can be restored by flipping other spins, at the cost of violating other ones, but without introducing additional energy (cf. Fig. 1.2). Using a chain of such spin flips, the constraints violated by the initial spin flip can be separated from one another, resulting in a *non-local* excitation of the same energy. Now, it is natural to break down these into a number of local excitations, associated with the violation of each constraint. These excitations are *fractionalised* on account of only being created in groups by physical excitations (spin flips): they carry fractional quantum numbers, and since they arise as parts of elementary excitations, they may have nontrivial topological properties and interactions.

Such a classical spin liquid can be endued with quantum dynamics by Hamiltonian terms that contain transverse spin components. These hybridise the ground-state manifold of the Ising model, leading to a ground state with massive *entanglement,* rather than degeneracy. The simplest approach is adding a weak transverse field; however, to lowest

[†]This degeneracy is *accidental,* not protected by any symmetry. In realistic systems, further-neighbour or long-range interactions always break the exact degeneracy; however, the temperature scale associated with this is often too low to matter in practice.

[‡]The ground states of the triangular lattice Ising model can be mapped onto dimer coverings of the honeycomb lattice formed by the centres of the triangles, allowing the exact entropy, $0.323\,k_\mathrm{B}$ per spin, to be calculated [8].

[§]The triangular-lattice Ising model fits this bill only partially: its ground-state manifold is constrained to $\sum \sigma = \pm 1/2$ on each triangle, which may or may not be violated by a spin flip.

order in perturbation theory, it only flips a single spin, which introduces fractionalised excitations that disrupt the ground-state manifold. Dynamics *within* this manifold is captured by higher-order *ring-exchange processes* acting on closed loops of spins so as to preserve all local constraints. This promotes the same constraints into conserved gauge charges, giving rise to *lattice gauge theories:* the ring-exchange dynamics generates fractionalised excitations of its own; the classical topological order is reflected in *topological quantum field theories,* which may manifest itself in anyonic quasiparticle statistics.

In the following sections, I shall discuss two clean-cut examples of ground-state degeneracy and fractionalisation in classical spin liquids, the classical six- and eight-vertex models [9], along with their quantum versions that give rise to U(1) and $\mathbb{Z}_2$ lattice gauge theories, respectively.

1.1.1 Classical six-vertex model

Let us first consider an antiferromagnetic Ising model defined on the *links* of a square lattice, where each link interacts with the six others it shares an endpoint with:

$$H = J \sum_{\langle ij \rangle} \sigma_i \sigma_j = \frac{J}{2} \sum_{+} \Big(\underbrace{\sum_{i \in +} \sigma_i}_{S_+} \Big)^2 - J \sum_i \sigma_i^2, \tag{1.1}$$

where $J > 0$ and the $\sigma_i = \pm 1/2$ are Ising variables; in the second form, the first summation runs over the vertices of the square lattice and $i \in +$ stands for the four links i that start in a given vertex. The original Ising Hamiltonian equals the cross-terms of the $(S_+)^2$ in the latter form; the subtracted σ_i^2 are irrelevant constants (namely, $1/4$). From the second form of (1.1), it is clear that any spin configuration in which $S_+ = \sum_{i \in +} \sigma^z$ is zero for all vertices is a ground state of the model: since there are $\binom{4}{2} = 6$ such configurations around each vertex (Fig. 1.1), Eq. (1.1) is called the *six-vertex model* [9, 10].

In what follows, it will be convenient to represent the Ising variables σ_i as vectors pointing from one end of the link it lives on to the other. As the square lattice is bipartite [that is, it can be separated into two (chequerboard) sublattices A and B with no links within a sublattice], we can choose to represent $\sigma = +1/2$ as a vector pointing from the A to the B sublattice, and $\sigma = -1/2$ as pointing from B to A; Fig. 1.1 shows this mapping for a vertex on the B sublattice. Now, the constraint that $S_+ = 0$ around each vertex requires that two vectors point into, as well as out of, each vertex.

It is easy to see that an extensive number of spin configurations satisfy this *two-in-two-out (2I2O) constraint.* Indeed, starting from one such configuration, we can choose any closed loop on the lattice such that the vectors on its links are joined start to end

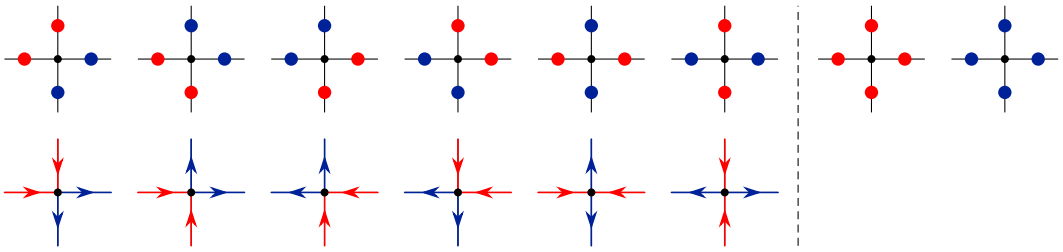

Figure 1.1. **Top row**: Spin configurations around a vertex allowed in the ground-state manifold of the six-vertex model (1.1) (left of dashed line) and the eight-vertex model (1.12) (all configurations). The red and blue dots stand for $\sigma = +1/2$ and $-1/2$, respectively. In the six-vertex model, each vertex is surrounded by two positive and two negative spins; in the eight-vertex model, all-positive and all-negative vertices are also allowed.
Bottom row: On a bipartite lattice, separated into two sublattices A and B, spin configurations can be represented as a lattice vector field: $\sigma = +1/2$ ($-1/2$) corresponds to a vector pointing from sublattice A to B (B to A). The figure shows this mapping for a vertex on sublattice B (see also Fig. 1.2). The six-vertex constraint is equivalent to a two-in-two-out (2I2O) rule on every vertex, which coarse-grains into a divergence-free constraint, allowing the six-vertex model to be treated as a U(1) Coulomb phase.

and flip it without violating the 2I2O rule: each vertex along the loop will only see one incoming and one outgoing vector swapped. One can build such a loop one vertex at a time, by following an outgoing link to a new vertex: as there are two outgoing links on each vertex to choose from, there is exponentially many of these loops, generating exponentially many ground states.

The number of these states can be estimated[†] through the famous argument of Pauling [10] for the residual entropy of water ice (§1.3.1): on a lattice of N vertices, there is a total of 2^{2N} spin configurations before taking into account any of the 2I2O rules; each of these rules eliminates 10 out of the 16 configurations possible around a given vertex. Assuming these rules are independent from one another, they reduce the number of available configurations by a factor of $(6/16)^N$, leading to the residual entropy

$$S = k_B \log \left[2^{2N} (6/16)^N \right] = N k_B \log(3/2). \tag{1.2}$$

Fractionalisation

Naïvely, the elementary excitation of any Ising model is flipping a spin: in the six-vertex model, this results in adjacent vertices with three-in-one-out (3I1O) and one-in-

[†]For the square lattice, the number of 2I2O configurations can be calculated exactly [11]: the exact entropy, $S = \frac{3}{2} N k_B \log(4/3)$, is about 6% larger than the Pauling estimate. The latter is, however, simpler, more instructive, and works to a similar accuracy in 3D, too.

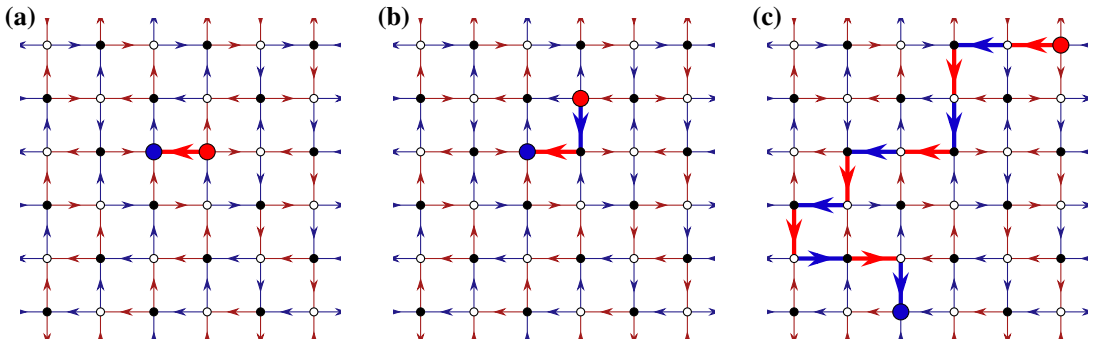

Figure 1.2. The elementary excitations of the six-vertex model (1.1) are single vertices that violate the 2I2O rule, called spinons.
(a) They are created in pairs of 3I1O (blue dot) and 1I3O (red dot) vertices by flipping individual spins (thick bright arrow).
(b) By flipping an other outgoing edge of the latter, that vertex can be restored to the 2I2O rule, but a 1I3O vertex is created at the other end of the same edge.
(c) Iterating this step allows for separating the two spinons to an arbitrary distance without disrupting the 2I2O rule anywhere between them: that is, spinons are deconfined, elementary excitations.
The colour of arrows indicates the sign of the Ising variables σ (red: $\sigma = +1/2$, blue: $\sigma = -1/2$); the direction of arrows is consistent with the lattice vector-field mapping $\sigma_i \hat{e}_i$; the chequerboard sublattices of the square lattice are indicated by the small dots (black: A, white: B). Note that the sign of σ alternates along the flipped path: this is necessary to enforce the 2I2O constraint $S_+ = 0$.

three-out (1I3O) spin configurations, both of which cost energy $J/2$ [Fig. 1.2(a)]. Now, flipping one of the other outgoing spins on the 1I3O vertex restores the 2I2O rule there, but makes the other endpoint of the same spin 1I3O [Fig. 1.2(b)]. By iterating this process [Fig. 1.2(c)], the 1I3O vertex can be moved independently of its 3I1O counterpart; likewise, the 3I1O vertex can be moved by flipping an incoming spin.

This means that the true elementary excitations of the six-vertex model are not spin flips, but rather 3I1O or 1I3O vertices, which are created in pairs by flipping a spin: these excitations are called *spinons*. Since they are created by splitting a single spin flip in two, they carry spin-1/2 quantum numbers, which is at odds with the integer angular momentum of physical excitations in spin systems: we say these spinons are *fractionalised;* they must be created in pairs to ensure that the overall excited state has physical quantum numbers, similar to excitations with fractional charge in fractional quantum Hall systems.

Coulomb phases

To better understand the behaviour of these fractionalised excitations, it is instructive to introduce a field-theoretical description. Indeed, the representation of spins as vectors turns the spin configurations into a lattice vector field: for now, we define the latter as an arrangement of vectors on the links of a lattice, which are parallel to the link they are on. In this language, the 2120 rule translates into a constraint of vanishing divergence on each vertex:

$$\mathrm{div}_r(\sigma_i \hat{e}_i) \equiv \sum_{r':\langle rr'\rangle} (\sigma_{rr'}\hat{e}_{rr'}) \cdot (\vec{r}' - \vec{r}) = \pm \sum_{i\in r} \sigma_i = 0. \tag{1.3}$$

Here, r and r' are neighbouring sites of the square lattice, and $\hat{e}_{rr'}$ is the unit vector along the bond that connects them, oriented from sublattice A to B: $\vec{r}' - \vec{r} = \pm\hat{e}_{rr'}$ depending on which sublattice site r belongs to, which gives rise to the sublattice-dependent sign in front of $S_r = \sum_{i\in r} \sigma_i$.

So far, this is an exact rewriting of the fine-grained 2120 rule. The lattice vector field can, however, be coarse-grained by averaging over a suitably large area around every point in space, resulting in the continuum flux field

$$\vec{B}(\vec{r}) = N^{-1} \sum_{i\sim\vec{r}} \sigma_i \hat{e}_i, \tag{1.4}$$

where the sum runs over the N links closest to $\vec{r}$. The lattice vectors $\vec{r}' - \vec{r}$ in (1.3) coarse-grain naturally into differentials; therefore, Eq. (1.3) itself translates into the continuum divergence-free condition $\nabla \cdot \vec{B} = 0$.

The energy of every divergence-free configuration of $\vec{B}(\vec{r})$ is equal; however, they are not equally likely. Since σ on any given link is zero after averaging over all spin configurations, $\langle\vec{B}\rangle = 0$; furthermore, a large value of $\vec{B}$ implies a large surplus of lattice vectors pointing in a particular direction, which can be realised in exponentially fewer ways than an even split. In a saddle-point approximation, these effects can be captured by the quadratic entropy functional [12, 13]

$$S = \text{const.} - \frac{\kappa}{2} \int \mathrm{d}^2 r |\vec{B}|^2; \tag{1.5}$$

that is, the free energy $-TS$ of the coarse-grained spin configurations has the same form as the energy of an electric or magnetic field. Together with the divergence-free constraint, this implies that correlations of $\vec{B}$ on a bipartite lattice have the same form as dipole–dipole interactions, decaying as $1/R^2$ in 2D. Due to the analogy between this de-

scription and electrostatics, we say that the six-vertex model realises a *Coulomb phase*.

Spinons break the 2I2O rule and the corresponding divergence-free constraint (1.3): in particular, 3I1O and 1I3O vertices have lattice divergences +1 and −1, respectively; owing to Gauss' theorem, this translates to $\nabla \cdot \vec{B} = \pm 1$ after coarse-graining. That is, spinons behave like quantised charges of the Coulomb phase (1.5), with entropic Coulomb interactions, which scale as $\log R$ in 2D.

Note, finally, that six-vertex models can be defined on higher-dimensional lattices with fourfold coordination, too; indeed, we shall discuss the case of the diamond/pyrochlore lattice in more detail in §1.3.1. By the same arguments, fractionalised spinons are the elementary excitations of these models, too; in a 3D Coulomb phase, correlations in $\vec{B}$ decay as $1/R^3$ and spinons interact via a $1/R$ entropic Coulomb interaction [12, 14].

1.1.2 Quantum six-vertex model: U(1) gauge theories

The simplest way of enduing the Ising model (1.1) with quantum dynamics is adding a transverse field in the x direction:

$$H = \frac{1}{2}\sum_{+}\Big(\sum_{i\in+}\sigma_i^z\Big)^2 - \varepsilon\sum_i \sigma_i^x = \frac{1}{2}\sum_{+}\big(S_+\big)^2 - \frac{\varepsilon}{2}\sum_i\big(\sigma_i^+ + \sigma_i^-\big), \tag{1.6}$$

where we set $J = 1$ for simplicity. Let us start with the case $0 < |\varepsilon| \ll 1$: the transverse perturbation introduces dynamics between the hitherto static Ising configurations, leading to hybridisation and energy splittings within the ground-state manifold. However, applying a single σ_i^x only flips one spin, which introduces a pair of oppositely charged spinons, thus moving out of the ground-state manifold. In higher orders of perturbation theory, it becomes possible to flip a closed loop of spins with alternating signs: in the vector language, these translate to vectors joined start to end [cf. Fig. 1.2(c)], so flipping them does not violate the 2I2O rule. The shortest nontrivial loop on a square lattice is a plaquette of four links: in fourth order of perturbation theory, we therefore obtain the terms

$$H_{\text{eff}} = -\frac{g}{2}\sum_{\square}\Big(\sigma_1^+\sigma_2^-\sigma_3^+\sigma_4^- + \text{H.c.}\Big) \tag{1.7}$$

that act within the ground-state manifold of the six-vertex model. In (1.7), $g = \mathcal{O}(\varepsilon^4)$ and the sum runs over the plaquettes of the lattice; $\sigma_{1,\dots,4}$ are the spins living on the four links around a given plaquette. Since any other process that preserves the 2I2O rules is at least sixth-order in perturbation theory, the full Hamiltonian (1.6) can be replaced with (1.7) as long as $\varepsilon \ll 1$.

Beyond the ground-state manifold, the ring-exchange Hamiltonian (1.7) preserves S^+

around each vertex in *any* spin configuration. That is, H_{eff} has an extensive number of integer conserved quantities, one for each vertex of the lattice, which give rise to a gauge symmetry: since S_+ commutes with H_{eff}, so does $\exp(i\chi S^z) = \prod_{i\in+} \exp(i\chi\sigma_i^z)$ for any χ. As the action of $\exp(i\chi\sigma_i^z)$ is to rotate σ_i around the σ^z axis by angle χ, this means that the four spins around each vertex can be rotated by the same angle without changing the physics: this gauge symmetry is termed U(1) because such rotations form a U(1) group parametrised by the rotation angle $0 \leq \chi < 2\pi$.

In fact, Eq. (1.7) can explicitly be rewritten as a pure U(1) lattice gauge theory [15]. To do so, we replace the σ^z with quantum rotor variables n that can take any half-integer value; instead of $\sigma^\pm$, we introduce ladder operators that link n to $n \pm 1$: these can be written as $e^{\pm i\phi}$, where the operators n and ϕ are canonically conjugate. Intuitively, $\sigma^\pm \sim e^{\pm i\phi}$ suggests that ϕ represents the XY angle of the spin; indeed, as conjugate to the half-integer n, it becomes an angular variable. To eliminate the unphysical values of n, we add a term proportional to n^2 to the Hamiltonian:

$$H = \frac{U}{2}\sum_i n_i^2 - g\sum_{\square} \cos\left(\phi_1 - \phi_2 + \phi_3 - \phi_4\right); \tag{1.8}$$

as $U \to \infty$, all $n \neq \pm 1/2$ states are projected out of the low-energy sector. To bring (1.8) to a more familiar form, let us return to the construction of lattice vector fields in §1.1.1; for convenience, however, the vectors will be replaced by variables defined on *directed* links of the square lattice, such that changing the direction of the link flips the sign of the field: this is equivalent to taking the projection of the vector variable onto the direction $\vec{r}' - \vec{r}$. In particular, we define the fields $e_{rr'}$ and $a_{rr'}$ as

$$e_{rr'} = \begin{cases} +n & \text{if } r \in A, r' \in B \\ -n & \text{if } r \in B, r' \in A, \end{cases} \qquad a_{rr'} = \begin{cases} +\phi & \text{if } r \in A, r' \in B \\ -\phi & \text{if } r \in B, r' \in A, \end{cases} \tag{1.9}$$

where A and B are the chequerboard sublattices of the bipartite square lattice. In this convention, the relationship between (1.8) and quantum electrodynamics becomes evident:

- The gauge symmetry of (1.7) can be written as $a_{rr'} \to a_{rr'} + \chi_r - \chi_{r'}$, where the χ_r are arbitrary for all vertices of the lattice, in clear analogy with the standard gauge transformation $\vec{A} \to \vec{A} + \nabla\chi$.

- The lattice divergence of e is the spinon charge of the Coulomb phase, which was promoted to a conserved gauge charge in (1.7).

- $\phi_1 - \phi_2 + \phi_3 - \phi_4$ is the sum of the a fields going around the plaquette, passing be-

tween the two sublattices at each link, that is, the lattice curl of a: in this notation,

$$H = \frac{U}{2}\sum_i e_i^2 - g\sum_{\square}\cos\operatorname{curl} a, \tag{1.10}$$

very similar to that of continuum quantum electrodynamics,

$$H = \int \mathrm{d}^d r \left(\frac{\varepsilon\vec{E}^2}{2} + \frac{\vec{B}^2}{2\mu}\right). \tag{1.11}$$

The key difference between the compact lattice theory (1.10) and (1.11) is the cosine (rather than quadratic) magnetic term that implies that curl a can be changed by 2π by instanton events without changing the physics. In two dimensions, these instantons are relevant in the renormalisation-group sense, leading to their proliferation that in turn leads to long-range order and gaps out the photon excitations [16]: this mechanism is discussed in great detail in §3.4.3. By contrast, instantons are irrelevant in the U(1) gauge theories borne out of 3D six-vertex models, so (1.10) can be coarse-grained into (1.11). These theories, therefore, describe a quantum-spin-liquid phase with gapless emergent photons, spinons that carry quantised electric charge, and Dirac magnetic monopoles: the physics of these quasiparticles is discussed in more detail in §2.1.2.

Note, finally, that (1.10) is only equivalent to the ring-exchange Hamiltonian (1.7) if $U \to \infty$,[†] an inconvenient limit to work in: in fact, a spin liquid can only arise because the variables n are half-integer, which makes the gauge theory *frustrated,* without a unique ground state; if the n were integers, the same limit would lead to a confined phase with $e \equiv 0$. In this unfrustrated theory, there is a phase transition between the large-U confined phase and a *deconfined* one, which coarse-grains into the Maxwell Hamiltonian (1.11) at a critical U/g [18]. In the latter phase, the discrete nature of e is irrelevant, so it is expected that the frustrated theory also has a deconfined Coulomb phase at low U [15]. The fate of this phase as $U \to \infty$ can only be studied numerically: on the diamond/pyrochlore lattice (§2.1.1), it was found to survive in the original ring exchange Hamiltonian [19].

[†]Using a large-S expansion, it is possible to derive (1.8) with $U \ll g$ [17], suggesting that the ring-exchange Hamiltonian (1.7) is always in a liquid phase as $S \to \infty$. In this approach, taking the $S \to 1/2$ limit is difficult, and a phase transition (potentially dependent on whether S is integer or half-integer) at intermediate S cannot be ruled out *a priori.*

1.1.3 Eight-vertex model: $\mathbb{Z}_2$ gauge theories

Let us now consider the Hamiltonian

$$H = -\sum_{+} \Big(\underbrace{\prod_{i\in+} \sigma_i^z}_{A_+} \Big) - \varepsilon \sum_i \sigma_i^x . \tag{1.12}$$

The first term again describes an Ising spin liquid; in its ground-state manifold, however, the σ^z have to multiply to $+1$,[†] rather than add up to 0. In addition to the spin configurations of the six-vertex model, this also includes states where all spins around a vertex are positive or negative (Fig. 1.1), hence (1.12) is called the *eight-vertex model* [9, 20]. This model also has extensive ground-state degeneracy (its ground-state entropy is $Nk_\mathrm{B} \log 2$) and similar fractionalised excitations to the spinons of the six-vertex model. However, since $A_+ = (-1)^{S_+}$, the integer spinon charge S_+ is replaced by a *spinon parity:* intuitively, merging two spinons with $S_+ = +1$ results in a vertex with $S_+ = +2$, i.e., four positive spins, which belongs to the eight-vertex ground-state manifold. Nevertheless, these $\mathbb{Z}_2$ spinons are still fractionalised quasiparticles, as flipping a single spin creates a pair of them, which can be separated from one another by further spin flips [Fig. 1.3(a)], similar to U(1) spinons in §1.1.1.

Likewise, the transverse field term of (1.12) still cannot act within the ground-state manifold directly, as acting with σ^x on a link introduces a pair of spinons at its ends. Therefore, a closed loop of spins must be flipped again to remain in the ground-state manifold. Unlike the six-vertex case, however, only the *parity* of the total σ^z around each vertex needs to be conserved, so any loop can be flipped, regardless of the sequence of spins along it. For $\varepsilon \ll 1$, therefore, the dynamics between ground states is dominated by the effective Hamiltonian

$$H_\mathrm{eff} = -\sum_{+} \Big(\underbrace{\prod_{i\in+} \sigma_i^z}_{A_+} \Big) - g \sum_{\square} \Big(\underbrace{\prod_{i\in\square} \sigma_i^x}_{B_\square} \Big); \qquad [g = O(\varepsilon^4)] \tag{1.13}$$

the product of σ^x around the plaquette stands for 16 perturbative processes that flip the plaquette regardless of its original configuration.

The dynamics of (1.13) can again be captured in terms of a lattice gauge theory. However, the conserved gauge charge is no longer the U(1) spinon charge S_+, but its *parity*, $A_+ = (-1)^{S_+}$. Accordingly, the local gauge transformations are restricted from the U(1)

[†]For this section, we switch to the convention that the eigenvalues of spin are ± 1, rather than $\pm 1/2$, to make handling large products of spin operators more convenient.

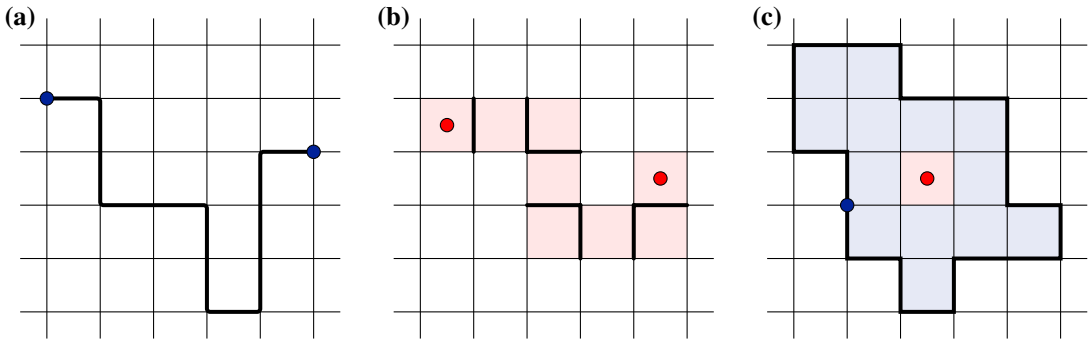

Figure 1.3. **(a)** Acting on the ground state with a σ^x operator flips A_+ at the two endpoints of the corresponding link, generating two spinons (blue dots). These can be separated at no additional energy cost by applying σ^x to consecutive links along a chain (black line).
(b) Acting on the ground state with a σ^z operator flips $B_\square$ on the two plaquettes adjacent to the corresponding link, generating two visons (red dots). Using further σ^z operators in the pattern of the black links, they can be separated at no extra cost (the corresponding path of the visons is shown in pink as a guide to the eye).
(c) A spinon (blue dot) can be carried around a closed loop (black line) using the product of σ^x on all links along the same. This operator can be rewritten as the product of plaquette operators $B_\square$ inside the loop, since links on the inside are included twice and thus contribute $(\sigma^x)^2 = 1$. If the loop encloses a vison (red dot), this product is -1, as all other plaquettes (light blue) contribute a factor of $+1$. That is, carrying a spinon around a vison flips the sign of the wave function, a hallmark of semionic mutual statistics.

group of the $e^{i\chi S_+}$ to the $\mathbb{Z}_2$ subgroup $\chi \in \{0, \pi\}$.

To better understand the resulting $\mathbb{Z}_2$ *gauge theory,* let us now consider the spectrum of the Hamiltonian (1.13), also known as the *toric code* [21]. Famously, all eigenstates of this model can be obtained exactly for arbitrary g, as not only do the $B_\square$ commute with the dominant Ising terms of (1.12), but *every term in the Hamiltonian commutes with every other:*

- the *star operators* A_+ and *plaquette operators* $B_\square$ commute among themselves since they are made up of commuting σ^z and σ^x operators, respectively;[†]
- the plaquette operators are designed to leave each star operator invariant, hence $[A_+, B_\square] = 0$ for each vertex and plaquette.

Accordingly, every A_+ and $B_\square$ also commutes with the Hamiltonian; that is, every eigenstate of (1.13) can be labelled with an eigenvalue of all star and plaquette operators. In

[†]The only non-commuting terms in the U(1) Hamiltonian (1.7), which make a similar treatment impossible, are the plaquette terms among themselves, as some contain σ^- and σ^+ for the same link.

a system with n vertices and periodic boundary conditions, each $A_+ = \pm 1$, $B_\square = \pm 1$ eigenvalue sets a condition to be satisfied by the eigenstate. However, as

$$\prod_+ A_+ = \prod_i \left(\sigma_i^z\right)^2 = +1; \qquad \prod_\square B_\square = \prod_i \left(\sigma_i^x\right)^2 = +1, \tag{1.14}$$

two of these constraints is redundant; the remaining $2n-2$ independent conditions define 2^{2n-2} sectors of degenerate eigenstates. As the Hilbert space of the system is 2^{2n}-dimensional (a total of $2n$ spins live on its links), each of these sectors contains four states, distinguished by global properties that cannot be picked up by local measurements. Before discussing the significance of these *topological degeneracies,* it is instructive to describe the spectrum in terms of elementary excitations.

In the ground state of (1.13) with $g > 0$, $A_+ = B_\square = +1$ for all vertices and plaquettes; excited states can be described in terms of quasiparticles that appear on vertices with $A_+ = -1$ and plaquettes with $B_\square = -1$, each carrying energy 2 and $2g$. The first are the $\mathbb{Z}_2$ spinons from the classical limit, also called *e particles*; the plaquette excitations are called *visons* or *m particles*. Due to (1.14), both species can only appear in pairs; that is, they are fractionalised. Spinons can be created on top of the ground state by acting with σ^x on some link: this leaves the $B_\square$ unchanged, but flips A_+ on the two ends of the same link, introducing a pair of spinons; applying further σ^x operators separate them at no extra cost [Fig. 1.3(a)]. Likewise, visons on adjacent plaquettes can be created and separated using σ^z operators [Fig. 1.3(b)]. Since σ^x operators commute among themselves, taking one spinon around another does not change the wave function, that is, they are bosons, and so are visons by the same argument. However, taking a spinon around a vison flips the sign of the wave function: $\prod \sigma^x$ around a loop is equal to the product of the $B_\square$ enclosed by it, which is -1 if one of these plaquettes hosts a vison [Fig. 1.3(c)]. Spinons and visons are, therefore, said to have *mutual semionic statistics:* an interesting consequence of this is that a bound state of a spinon and a vison (two bosons), called the ε particle, is a fermion![†]

Now, consider creating a pair of spinons in an eigenstate, carrying one around a loop that winds around the periodic boundary conditions, and recombining them: clearly, doing so does not change any of the state labels $\{A_+, B_\square\}$, but as $\prod \sigma^x$ around this loop is not a combination of the A_+ and $B_\square$, the final state is generally different from the

[†]The situation is altogether similar to the physics of superconductors, where the condensation of Cooper pairs breaks electric charge conservation [a U(1) symmetry] down to a $\mathbb{Z}_2$ parity symmetry [22]. This leads to the expulsion of the original U(1) gauge field (electromagnetism), except for vortices with flux $h/(2e)$. The gapped excitations of a superconductor are the same vortices and Boguliubov quasiparticles with charge $\pm e$: the vortices act as trivial 2π-fluxes for Cooper pairs, but as π-fluxes for the Boguliubov quasiparticles, leading to semionic mutual statistics with the latter.

initial one. There are two independent loop operators of this type, one each for loops winding around horizontal and vertical boundary conditions: since they commute with one another and all star and plaquette operators, they can be used to label the four degenerate eigenstates in each $\{A_+, B_\square\}$ sector in a physically meaningful manner.[†]

The fact that these four states are only distinguished by global observables (such as the loop operators), but not local measurements like A_+ or $B_\square$, is a hallmark of *topological order*.[‡] In particular, $\mathbb{Z}_2$ topological order is characterised by global $\mathbb{Z}_2$ topological invariants and a 2^d-fold *topological ground-state degeneracy* in d dimensions. In 2D, it also gives rise to a constant universal correction of $-\log 2$ to the area law of entanglement entropy [23–25].

$\mathbb{Z}_2$ gauge theories are appealing platforms for quantum computing on account of their topological degeneracy [21]. Namely, the four degenerate ground states of (1.13) can be used to represent two qubits: the exact degeneracy stops any dephasing between qubit components and the quantum information is not perturbed by incoherent local noise, since the ground states are indistinguishable to local operators. The toric code was in fact designed as an idealised $\mathbb{Z}_2$ gauge theory to demonstrate these properties [21]. Nevertheless, the lack of continuous gauge fields amenable to coarse-graining in more realistic theories [e.g., (1.12), or the dimer models of §3.1.3] still leads to short-range correlations with some finite correlation length ξ [18, 26]; accordingly, ground states in different topological sectors remain degenerate and indistinguishable to local probes up to $O(e^{-L/\xi})$ differences, where L is the linear system size.

[†]Similar loop operators involving σ^z spin components can be defined using visons. These, however, anticommute with the σ^x loop operators, so acting with them flips the eigenvalues of the same.

[‡]A form of topological order arises in U(1) gauge theories as well, since the flux of $\vec{e}$ across any surface that spans the entire system can only be changed by loop updates that wind around the periodic boundary conditions and intersect the same surface [15]. The resulting topological sectors, however, can be distinguished by local measurements, as the topological invariant gives rise to a background $\vec{e}$-field, as discussed in §3.1.2. There is no topological correction to entanglement entropy, either.

1.2 Heisenberg antiferromagnets

> *To every thing there is a season, and a time to every purpose under the heaven: a time to break down, and a time to build up; a time to cast away stones, and a time to gather stones together; a time to rend, and a time to sew.*
>
> — from Ecclesiastes 3:1–8

The perturbative construction in the previous section provides transparent access to the topological orders and gauge theories that underpin QSL behaviour. Several paradigmatic frustrated magnets (including the J_1–J_2 model discussed in chapter 4) are, however, better described through isotropic Heisenberg interactions:

$$H = \sum_{ij} J_{ij} \vec{s}_i \cdot \vec{s}_j, \tag{1.15}$$

where the $\vec{s}_i$ are spin-1/2 angular momentum operators on lattice sites.[†] Due to isotropy, no component of the interaction can be designated as a dominant classical component, undermining the construction above; doing so also breaks the SU(2) spin-rotation symmetry of the Hamiltonian explicitly, which should be preserved by the spin-liquid ground state.

Historically, the first approach to describing spin liquids in Heisenberg antiferromagnets (and QSLs in general) was Anderson's *resonating-valence-bond (RVB)* ansatz [27–30], in which nearby pairs of spins are combined into singlets: the RVB state arises as a superposition of all possible singlet coverings. The spin-liquid properties of RVB states, their fractionalised excitations, and the quantum dimer models that capture their quantum dynamics are discussed in detail in chapter 3.

An alternative route [29, 31] to capturing SU(2)-symmetric spin liquids starts from the observation that Heisenberg models can be obtained as the limit of the Fermi–Hubbard model

$$H = -t \sum_{\langle ij\rangle,\sigma=\uparrow,\downarrow} \left(c^\dagger_{i\sigma} c_{j\sigma} + c^\dagger_{j\sigma} c_{i\sigma} \right) + U \sum_i n_{i\uparrow} n_{i\downarrow} \tag{1.16}$$

at half-filling, as $U \to \infty$. In this limit, the Hubbard interaction forces every site to be singly occupied, so the hopping term only acts perturbatively, by exchanging two electrons: owing to the SU(2) symmetry of (1.16), the effective exchange interaction will

[†] In this section, we denote spin operators with s rather than σ to avoid confusing with Pauli matrices that feature prominently in SU(2) parton constructions.

take the Heisenberg form (1.15) with $J = 4t^2/U$ between nearest neighbours [32]. The spin operators are composed out of the electron operators c as (summation implied over the Greek spin indices)

$$\vec{s}_i = \frac{1}{2} c^\dagger_{i\alpha} \vec{\sigma}_{\alpha\beta} c_{i\beta}, \tag{1.17}$$

where $\vec{\sigma}$ is the vector of Pauli matrices; it is straightforward to verify that these operators satisfy angular momentum commutation relations.

We can now reverse this argument by replacing the spin operators in any spin model with a pair of *Abrikosov fermions*, $c_{\uparrow,\downarrow}$, defined by (1.17). This is a faithful representation of the SU(2) algebra as long as each site is occupied by precisely one fermion:

$$c^\dagger_\uparrow c_\uparrow + c^\dagger_\downarrow c_\downarrow = 1; \qquad c^\dagger_\uparrow c^\dagger_\downarrow = 0; \qquad c_\uparrow c_\downarrow = 0; \tag{1.18}$$

in general, however, the original spins are not borne out of a pair of electron orbitals, so the Abrikosov fermions are fictitious quasiparticles. Now, the Heisenberg Hamiltonian (1.15) is replaced by the interaction-only fermionic model

$$H = \sum_{ij} J_{ij} \left(\frac{1}{2} c^\dagger_{i\alpha} c^\dagger_{j\beta} c_{j\alpha} c_{i\beta} - \frac{1}{4} \right). \tag{1.19}$$

In the first approximation, such Hamiltonians can be treated in mean-field theory, by performing an appropriate Hubbard–Stratonovich transformation on the four-fermion terms of (1.19). To describe a spin liquid, we require that the mean-field amplitudes in this decoupling be invariant under simultaneous SU(2) rotations of all spins: this excludes, among others, the Curie–Weiss approximation $\vec{s}_i \cdot \vec{s}_j \longrightarrow \langle \vec{s}_i \rangle \cdot \vec{s}_j + \vec{s}_i \cdot \langle \vec{s}_j \rangle - \langle \vec{s}_i \rangle \cdot \langle \vec{s}_j \rangle$ used to describe ordered magnets. In general, there are two appropriate decoupling channels [31, 33, 34],

$$\frac{J_{ij}}{2} c^\dagger_{i\alpha} c^\dagger_{j\beta} c_{j\alpha} c_{i\beta} \longrightarrow - \underbrace{\frac{J_{ij}}{2} \left\langle c^\dagger_{i\alpha} c_{j\alpha} \right\rangle}_{t^*_{ij}} c^\dagger_{j\beta} c_{i\beta} - \underbrace{\frac{J_{ij}}{2} \left\langle c^\dagger_{j\beta} c_{i\beta} \right\rangle}_{t_{ij}} c^\dagger_{i\alpha} c_{j\alpha} + \text{const.}; \tag{1.20a}$$

$$\frac{J_{ij}}{2} c^\dagger_{i\alpha} c^\dagger_{j\beta} c_{j\alpha} c_{i\beta} \longrightarrow \underbrace{\frac{J_{ij}}{2} \left\langle \varepsilon_{\alpha\beta} c^\dagger_{i\alpha} c^\dagger_{j\beta} \right\rangle}_{\Delta_{ij}} \varepsilon_{\gamma\delta} c_{j\gamma} c_{i\delta} + \underbrace{\frac{J_{ij}}{2} \left\langle \varepsilon_{\alpha\beta} c_{j\alpha} c_{i\beta} \right\rangle}_{\Delta^*_{ij}} \varepsilon_{\gamma\delta} c^\dagger_{i\gamma} c^\dagger_{j\delta} + \text{const.}, \tag{1.20b}$$

resulting in the mean-field Hamiltonian

$$H_{\text{MF}} = \sum_{ij} \left(-t_{ij} c^\dagger_{i\alpha} c_{j\alpha} + \Delta^*_{ij} \varepsilon_{\alpha\beta} c^\dagger_{i\alpha} c^\dagger_{j\beta} + \text{H.c.} \right) + \sum_i \lambda_i c^\dagger_{i\alpha} c_{i\alpha}; \tag{1.21}$$

the last term is a Lagrange multiplier to ensure that each site has average occupancy 1.

The elementary excitations of the quadratic fermion Hamiltonian (1.21) are Bogoliubov quasiparticles made out of c and $c^\dagger$ operators. These excitations are again fractionalised, for they carry spin-1/2 quantum numbers [this time, SU(2) spin rather than s^z component], and the physical spin operators (1.17) create them in pairs. They are called *spinons* in analogy with those discussed in §1.1; the name *parton* is also used as the Abrikosov fermions are "parts" of the physical spin degrees of freedom. We note that similar mean-field theories can also be constructed starting from the Schwinger-boson representation

$$\vec{s}_i = \frac{1}{2} b^\dagger_{i\alpha} \vec{\sigma}_{\alpha\beta} b_{i\beta} \qquad \left(b^\dagger_{i\alpha} b_{i\alpha} = 2S\right) \tag{1.22}$$

of spin operators $\vec{s}$ of arbitrary SU(2) quantum number S; indeed, the Heisenberg Hamiltonian (1.19) works out identically to the fermionic case. An advantage of the bosonic formalism is that the mean-field theory becomes exact in the limit of large boson numbers (that is, large S), allowing a controlled approximation of the spin-1/2 limit through a $1/S$ expansion.

Consider now a mean-field Hamiltonian (1.21) in which all Δ terms vanish: such a Hamiltonian is invariant under rotating the phases of all c operators, $c_{i\sigma} \to e^{i\chi} c_{i\sigma}$ [33]. This symmetry can be promoted to a U(1) gauge symmetry by introducing the lattice vector potential A_{ij} as the phase of the spinon hopping terms [35]:

$$H_{\mathrm{U}(1)} = \sum_{ij} \left(t_{ij} e^{iA_{ij}} c^\dagger_{i\alpha} c_{j\beta} + \mathrm{H.c.} \right) + H_\mathrm{g}, \tag{1.23}$$

where H_g only contains gauge-invariant combinations of the vector potential A and the dual "electric" field E, cf. §1.1.2.[†] Equation (1.23) is invariant under the U(1) gauge transformation

$$A_{ij} \to A_{ij} + \chi_i - \chi_j \qquad c_{i\alpha} \to e^{i\chi_i} c_{i\alpha} \qquad c^\dagger_{i\alpha} \to e^{-i\chi_i} c^\dagger_{i\alpha};$$

spinons become (fermionic) matter minimally coupled to a compact U(1) gauge theory. The behaviour of such systems can be qualitatively different from pure gauge theories [7]; in particular, spinon Fermi surfaces or Dirac fermions can stabilise a U(1) QSL even in two dimensions.

In the presence of pairing terms Δ, the U(1) phase-rotation symmetry breaks down to

[†]Pure gauge contributions also arise from the ground-state energy of the first term of (1.23), which clearly depends on A, but only in a gauge-invariant fashion.

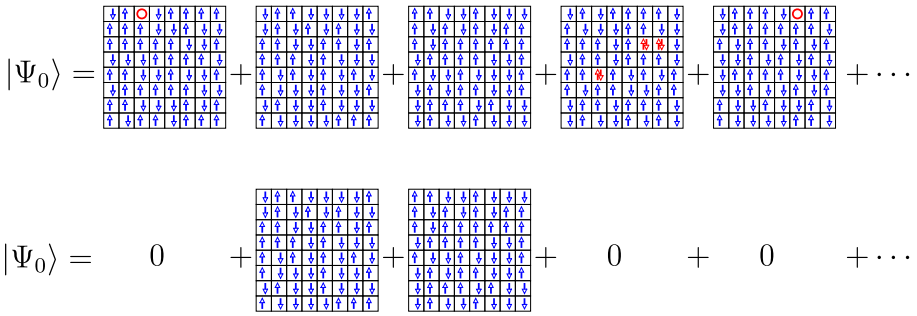

Figure 1.4. The Gutzwiller projection. The mean-field wave function $|\Psi\rangle$ is expanded in the basis of fermion number eigenstates; then, all basis states with empty or doubly occupied sites are removed and the rest is interpreted as s^z basis states. In practical algorithms, such as variational Monte Carlo, the projection can be implemented by not sampling basis states that do not correspond to a s^z basis state [38]. Figure taken from Ref. 7.

the $\mathbb{Z}_2$ subgroup $\chi = \{0, \pi\}$; accordingly, spinons in such mean-field theories are coupled to a $\mathbb{Z}_2$ gauge theory. Due to their fermionic statistics, these spinons are equivalent to the ε particles of the toric code; bosonic spinons (e particles) can be constructed analogously from Schwinger-boson mean-field theories. In the $\mathbb{Z}_2$ gauge theory, however, one can be transmuted into the other by binding a vison to it: this illustrates the commonly observed equivalence of fermionic and bosonic parton approaches.

In general, QSL mean-field Hamiltonians on a given lattice can exhibit a variety of symmetry properties and emergent gauge theories. These can be classified using *projective symmetry group* analysis [36, 37], which supplants Landau theory for QSL phases that cannot be characterised in terms of broken symmetries. The classification imposes symmetry constraints on the available mean-field Hamiltonian terms: this is important in quantitative approaches and may constrain spinon dispersions qualitatively (e.g., through symmetry-protected gapless points).

Beyond the qualitative description of QSLs in terms of spinons coupled to lattice gauge theories, parton mean-field theories can be used quantitatively in several ways:

- Equation (1.21) can be solved self-consistently, by enforcing the definition (1.20) of t and Δ, and setting λ on each site so as to ensure $\langle c_\alpha^\dagger c_\alpha \rangle = 1$. This is an uncontrolled approximation due to the mean-field decoupling of (1.19) and the relaxation of the hard constraint (1.18) on fermion numbers; it can be made more rigorous in the bosonic case using a $1/S$ expansion.

- Following Gutzwiller's construction [39] of variational ansätze for the Hubbard model (1.16), one can take the ground-state wave function of the noninteracting Hamiltonian (1.21), and project out all number basis states that are inconsistent

with (1.18), see Fig. 1.4. The construction of the Heisenberg model as the limit of Hubbard models implies that such *Gutzwiller-projected mean-field wave functions* capture qualitative features of the true ground state accurately, providing a powerful ansatz for variational approaches [29, 31]. The coefficients t, Δ, λ of the Hamiltonian are not found self-consistently, but are treated as variational parameters to minimise the energy of the variational state [38].

- The Hubbard–Stratonovich decoupling (1.21) is exactly equivalent to (1.19) provided t, Δ, and the Lagrange multipliers λ are treated as dynamical quantum fields. Starting from the mean-field values, perturbative techniques (such as the functional renormalisation group) allow these fluctuations to be absorbed in the c operators, resulting in exact quasiparticles whose spectrum is captured quantitatively by a renormalised quadratic Hamiltonian.

1.3 Rare-earth pyrochlores

> *A complicated structure? Undoubtedly. But after all, the cathedral of Milan is complicated too, and you still look at it with awe.*
>
> — Kató Lomb

In this section, I discuss materials of the empirical formula $A_2B_2O_7$, where A is a 3+ rare-earth cation and B is a 4+ cation of transition metals or IV.B elements: this class includes a great variety of experimentally and theoretically important frustrated magnets, including classical and quantum spin ice (chapter 2) and $Ho_2Ir_2O_7$ (chapter 5). This is because $A_2B_2O_7$ crystallises in the *pyrochlore lattice* for a wide range of elements A and B (Fig. 1.5). This lattice belongs to the space group $Fd\bar{3}m$ in the FCC crystal system; each cubic unit cell contains eight formula units. The A and B metal ions live on the vertices of two interpenetrating networks of *corner-sharing tetrahedra* [Fig. 1.6(a)], which can be decomposed into four FCC sublattices each. These networks (also commonly called pyrochlore lattices) can be constructed as the link midpoints of two diamond lattices, shifted from one another by $\frac{1}{2}\langle 100\rangle$ [Fig. 1.6(b)]. Eight oxygen atoms per unit cell (labelled O') live in the centres of A atom tetrahedra; the remaining 48 atoms (labelled O) occupy a set of $48f$ Wyckoff positions. A and B ions have cubic and octahedral oxygen coordination, respectively [Fig 1.6(d,e)]; due to the low symmetry of the $48f$ sites [40], the coordination polyhedra are distorted and can be continuously deformed by (chemical) pressure.

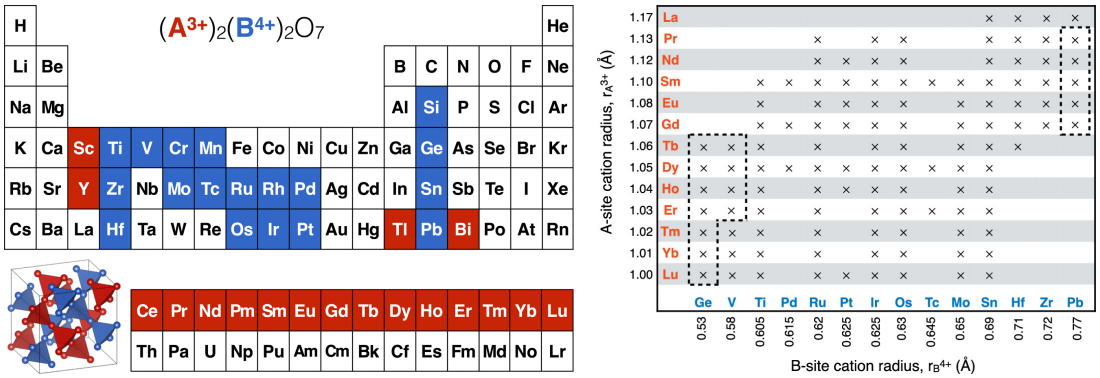

Figure 1.5. Summary of all known pyrochlores of cation valences 3+ and 4+.
Left: Elements that appear in pyrochlores with valence 3+ (red) or 4+ (blue), shown in the periodic table. Bottom left corner: intercalated structure of A and B site cations.
Right: Table of stable rare-earth pyrochlores. Species inside the dashed lines are prepared under applied pressure; those outside are stable at ambient pressure.
Figures taken from Ref. 41.

La	Ce	Pr	Nd	Pm	Sm	Eu	Gd	Tb	Dy	Ho	Er	Tm	Yb	Lu
0	5/2	4	9/2	4	5/2	0	7/2	6	15/2	8	15/2	6	7/2	0

Table 1.1. Total angular momentum J of 3+ rare-earth cations [44].

The pyrochlore lattice can also be constructed by stacking alternating kagome and triangular layers along a $\langle 111 \rangle$ direction. This shows that shortest nontrivial loops around which ring-exchange interactions are possible are the hexagonal plaquettes of the kagome layers. B ions are located in the centres of such plaquettes formed out of A ions and vice versa [Fig. 1.6(c)]: this duality between pyrochlore lattices will be important both for quantum spin ice (chapter 2) and pyrochlore iridates (chapter 5).

The net angular momentum of 3+ rare-earth ions reaches quite substantial values (Table 1.1); the corresponding large spin degeneracy is, however, lifted by crystal electric field (CEF) effects due to their oxygen environment. The ground state of the single-ion CEF Hamiltonian is usually a doublet [45–47]: for Kramers (half-integer spin) ions, it is protected by time-reversal symmetry [48, 49]; for non-Kramers (integer spin) ions, it is due to lattice symmetry, and so it can be lifted by disorder or other lattice distortions. This doublet is typically separated from higher CEF states by a gap on the order of hundreds of kelvins [45–47], much greater than the energy scale of superexchange interactions between the rare-earth moments (usually a few kelvins); as a result, low-temperature physics is captured accurately by the ground-state doublet only, allowing for effective spin-1/2 descriptions.

The properties of the ground-state doublet are determined by CEF effects, which in

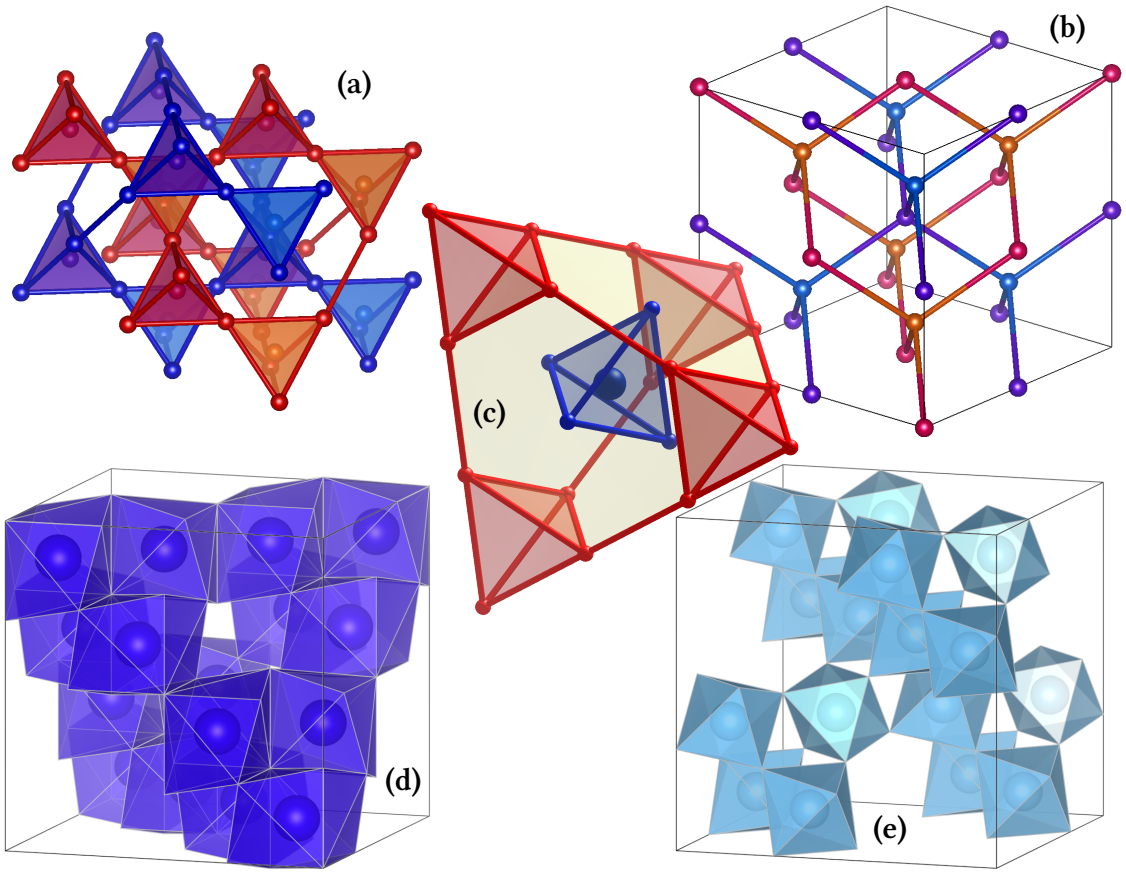

Figure 1.6. The pyrochlore crystal structure.
(a) The A^{3+} and B^{4+} metal ions sit on the sites of intercalated pyrochlore lattices (shades of blue and red, respectively), shifted from one another by $\frac{1}{2}\langle 100\rangle$. Each lattice is made up of corner-sharing tetrahedra, the vertices of which host the ions. The centres of these tetrahedra form two diamond lattices **(b)**; the pyrochlore sites are the link midpoints of the latter. The diamond lattice is bipartite (different shades of blue and red); the two sublattices correspond to translation-inequivalent "up" and "down" tetrahedra of the pyrochlore lattices.
(c) The shortest nontrivial loops of a pyrochlore lattice (red) are (regular) hexagonal plaquettes that lie in {111} planes (yellow). The midpoints of these plaquettes (blue) form another, dual pyrochlore lattice; the dual of this lattice is, in turn, the original one. The pyrochlore lattices formed by A and B ions are the duals of one another.
(d) Oxygen coordination polyhedra of A^{3+} ions. Each ion is surrounded by two O' ions in the centres of the tetrahedra it belongs to and six O ions at $48f$ Wyckoff positions, whose positions depend on a single parameter x [40]. These eight ions form a distorted cube (perfect cube for $x = 3/8$) around each A ion; the cubes around neighbouring ions share an edge, which meet on a single O' ion for all four A ions on a tetrahedron. Figure based on the structure of $Dy_2Ti_2O_7$ ($x = 0.3214$ [42]); oxygens not shown to reduce clutter.
(e) Oxygen coordination polyhedra of B^{4+} ions. Each ion is surrounded by six O ions that form a distorted octahedron (perfect octahedron for $x = 5/16$); the octahedra around neighbouring B ions share a single oxygen ion.
Figures created using VESTA 3 [43].

turn depend strongly on the exact oxygen geometry around the rare-earth ions. Likewise, magnetic interactions between the same moments is dominated by superexchange, whose intensity is very sensitive to the position of oxygen ions, as well as the structure of the rare-earth *f*-orbitals. As a result, subtle changes in chemical pressure allow rare-earth pyrochlores to realise a wide variety of effective spin-1/2 Hamiltonians between easy-axis, easy-plane, and Heisenberg-type spins [50, 51], making them a popular experimental platform for realising frustrated magnetic phases. In the rest of this section, I shall focus on *spin ice* [52, 53], an archetypal frustrated Ising model realised most cleanly by the pyrochlore titanates $Dy_2Ti_2O_7$ and $Ho_2Ti_2O_7$: beyond their significance as the best established experimental examples of spin-liquid physics, our excellent theoretical understanding of spin ice serves as the basis of discussing other pyrochlore systems. Later in this dissertation, I shall return to *quantum spin ice* (chapter 2), as well as *pyrochlore iridates* (chapter 5) in which the magnetism of Ir^{4+} ions introduces qualitative changes in behaviour compared to spin ices with a non-magnetic *B* site.

1.3.1 Nearest-neighbour spin ice

In §1.1, extensive ground-state entropy was identified as a hallmark of classical spin-liquid phases. The first system where such residual entropy was observed was, however, not a spin system but water ice, in which thermodynamic measurements in the 1930s [54, 55] indicated substantial entropy remaining down to the lowest temperatures. This can be understood in terms of the crystal structure of ice: oxygen atoms form a regular lattice[†] with fourfold coordination; hydrogen atoms live along the links of this lattice, bound to one oxygen atom by a hydrogen bond, and covalently to the other. Due to the difference in strength between these bonds, the hydrogen sits closer to the covalently bound oxygen: there are two equilibrium hydrogen positions on each O–O link, one closer to each end, and four near each oxygen atom. The arrangement of hydrogen atoms on these sites is constrained by the Bernal–Fowler *ice rules* [56]:

(i) each O–O link hosts precisely one hydrogen atom;

(ii) each oxygen atom has two covalent and two hydrogen O–H bonds.

Due to rule (i), each O–O link is characterised by which of the two equilibrium positions is occupied by a hydrogen atom: we can represent this information with arrows on the links that point towards the covalently bound oxygen. Now, rule (ii) requires that two

[†]In the metastable I_c phase, the oxygen atoms sit on diamond lattice sites; in the more common I_h phase, they form a hexagonal lattice that can be derived from the diamond lattice by changing the stacking of its {111} layers, similar to the difference between HCP and CCP packing.

of these arrows point towards, and away from, each oxygen atom: this is exactly the 2I2O constraint that defines six-vertex models (§1.1.1); in fact, the first such model was proposed by Pauling [10] to obtain the estimate (1.2) of the residual entropy of ice.

As seen in §1.1.1, the same six-vertex model is captured by the ground-state manifold of the antiferromagnetic Ising model

$$H = J \sum_{\langle ij \rangle} \sigma_i \sigma_j = \frac{J}{2} \sum_t \Big(\sum_{i \in t} \sigma_i \Big)^2 + \text{const.}; \tag{1.24}$$

due to the coordination of ice, the link midpoints i around each vertex form tetrahedra, labelled by t; the antiferromagnetic interactions are now really between nearest-neighbour spins. This Hamiltonian was first written down for the pyrochlore lattice[†] by Anderson [57] to discuss the structural disorder of inverse spinel materials.

As an actual spin model, Eq. (1.24) is unrealistic, as there is no single Ising axis that would be consistent with the cubic symmetry of the pyrochlore lattice. However, the $\langle 111 \rangle$ direction of the diamond-lattice link each rare-earth atom sits on is unique under the same symmetries, marking it out as a candidate (sublattice-dependent) Ising axis; from now on, the Ising variables σ and the z direction shall refer to these *local* $\langle 111 \rangle$ *axes.* Indeed, it was discovered that the CEF ground-state doublets of $Ho_2Ti_2O_7$ [45, 58] and $Dy_2Ti_2O_7$ [46] have very strong easy-axis anisotropy along these directions: they are predominantly made up of the maximally z-polarised states $m_J = \pm 8$ and $m_J = \pm 15/2$, respectively, so the g-tensor that couples the effective spin-1/2 operators to external magnetic fields has a g^{zz} component around 20, orders of magnitude greater than the transverse components.[‡] That is, the angular momentum operator $\vec{J}_i$ can effectively be replaced by $m_J^{\max} \sigma_i \hat{z}_i$, where $m_J^{\max} = 8$ for Ho, and 15/2 for Dy (Table 1.1), the σ_i are Ising variables, and $\hat{z}_i$ is a unit vector along the local $\langle 111 \rangle$ axis. Owing to the analogy with Pauling's ice model, such systems are called *spin ices* [58]; as expected, they exhibit a residual entropy consistent with that of ice [59].

Projecting any nearest-neighbour interaction onto this ground-state doublet will yield an Ising model of the form (1.24); however, the fact that local $\langle 111 \rangle$ axes of neighbouring sites make a 109.5° angle with one another leads to a counterintuitive relationship between the original interactions and the projected Ising model. For example, the projection of a Heisenberg coupling, $J \vec{\sigma}_i \cdot \vec{\sigma}_j$ is $-J \sigma_i \sigma_j / 3$: to end up with an antiferromagnetic Ising model (1.24), one has to start with *ferromagnetic* interactions, which reflects in

[†] Since the pyrochlore lattice consists of the links of a diamond lattice [Fig. 1.6(b)], the spin-ice Hamiltonian (1.24) also captures the structural disorder of ice I_c exactly.

[‡] Due to the threefold rotation symmetry of the oxygen coordination around the local $\langle 111 \rangle$ axis, the only allowed anisotropy of such a g tensor is g^{zz} being different from $g^{xx} = g^{yy}$.

a positive Curie–Weiss temperature, unusual for a frustrated magnet [58]. Indeed, 2120 configurations maximise the net magnetic moment of each tetrahedron ($4\mu/\sqrt{3}$ along one of the $\langle 100 \rangle$ axes). By contrast, antiferromagnetic interactions generate a ferromagnetic Ising model with $J < 0$ in (1.24): all σ are equal in the ground state of this Hamiltonian, which translates to all spins pointing from "down" to "up" tetrahedra or vice versa; as expected for an antiferromagnet, these *all-in-all-out (AIAO)* ordered states have no net magnetic moment on any tetrahedron.

Pinch points

Since (1.24) is identical to the classical six-vertex model of §1.1.1 on the diamond lattice, its key features are the same too; furthermore, the vector mapping (Fig. 1.1) of the Ising degrees of freedom is particularly natural, as the physical spins $\vec{J}$ are actually parallel to the diamond-lattice links they are on. In particular, spin ice has an extensive ground-state degeneracy of 2120 states, with a low-temperature entropy given by (1.2). Spin–spin correlations within this ground-state sector can be extracted from the 3D Coulomb phase partition function (1.5): in particular, the correlators of the coarse-grained flux field $\vec{B}$ are given by [12]

$$\left\langle B_\mu(\vec{0}) B_\nu(\vec{r}) \right\rangle = \frac{4\pi}{\kappa} \left[\delta(\vec{r}) + \frac{1}{r^3} \left(\delta_{\mu\nu} - \frac{3 r_\mu r_\nu}{r^2} \right) \right], \tag{1.25}$$

similar to the functional form of dipolar interactions, decaying algebraically at long distances. In practice, it is easier to deal with reciprocal space correlators: Fourier transforming (1.25) gives

$$\left\langle B_\mu(-\vec{k}) B_\nu(\vec{k}) \right\rangle = \frac{1}{\kappa} \left(\delta_{\mu\nu} - \frac{k_\mu k_\nu}{k^2} \right) \tag{1.26}$$

in the long-wavelength limit (i.e., $k \ll a_0^{-1}$). Importantly, the correlator only depends on the direction of $\vec{k}$ but not on its magnitude: this results in a (non-divergent) singularity at the Γ point, which is a hallmark of a long-range correlated critical state.

Experiments are generally not able to measure the full correlation tensor (1.26), but only a projection of it selected by the physics of the probe itself. The singularity of (1.26) at the Γ point translates into singular features in the measured correlators. For example, the correlations of σ^z near the (002) reciprocal lattice vector are given by [12, 60]

$$\left\langle \sigma\left[(002) + \vec{k}\right] \sigma\left[(00\bar{2}) - \vec{k}\right] \right\rangle \propto \frac{k_x^2 + k_y^2}{k^2}, \tag{1.27}$$

which shows up as a bow-tie like feature in an intensity plot [61]: these, commonly called *pinch points,* are the most important experimental signature of spin ices (§1.3.3).

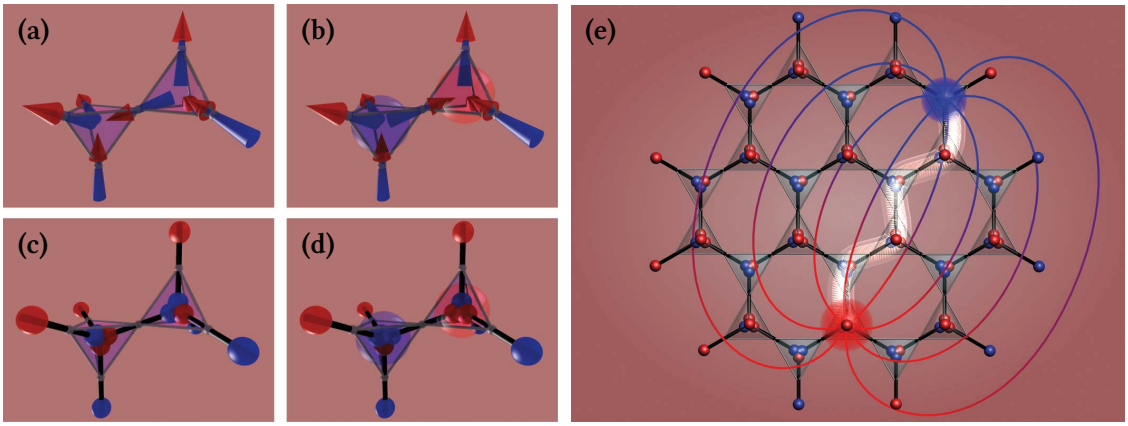

Figure 1.7. Fractionalisation in spin ice and the dumbbell model.
(a) Both tetrahedra satisfy the 2120 ice rules.
(b) By flipping one spin, one tetrahedron gets three incoming, the other three outgoing spins, hence developing a "positive" and "negative" spinon.
(c) The spins are replaced with dumbbells of a positive and a negative magnetic monopole, located at the sites of the diamond lattice. This introduces quadrupolar corrections which can be ignored safely.
(d) As a result, spinons develop an effective magnetic charge $\pm 2\mu/a_{\mathrm{d}}$.
(e) Spinons interact via an entropic (in nearest-neighbour spin ice) or an energetic (in dipolar spin ice) Coulomb force; they are not confined. They do not satisfy Dirac quantisation, but the choice of the Dirac string (white) is not unique and has little physical importance.
Figure taken from Ref. 62.

Spinons

Being a the six-vertex model, the elementary excitations of spin ice are fractionalised spinons, that is, 3110 and 1130 tetrahedra,[†] created in pairs (Fig. 1.7). As discussed in §1.1.1, they are positive and negative charges in the 3D Coulomb phase (1.5); accordingly, they interact through a $1/R$ entropic Coulomb interaction.

Creating spinons has a finite energy cost $\Delta = J/2$ (on the order of a few kelvins in $Dy_2Ti_2O_7$ and $Ho_2Ti_2O_7$); accordingly, their density depends on temperature as $\rho \propto e^{-\Delta/T}$. These thermally activated spinons can be treated as a 3D Coulomb gas of interacting charged particles: such gases are always in the Debye plasma phase [63], where their Coulomb interactions are screened by the surrounding charges over the Debye length $\xi \propto \sqrt{\rho}$ [64], leading to spinon–spinon correlations of the form

$$\langle q(0)q(\vec{r})\rangle \sim \frac{e^{-r/\xi}}{r} \implies \langle q(-k)q(k)\rangle \sim \frac{1}{1+\xi^2k^2}, \tag{1.28}$$

cf. Appendix A.3.1.

[†]In the context of water ice, these defects correspond to H_3O^+ and OH^- ions.

The origin of the dipolar correlations (1.26) and the corresponding pinch points in spinon-free spin ice is the divergence-free constraint $\nabla \cdot \vec{B} = 0$, which is violated by the spinons. It follows that the originally long-range algebraic spin correlations become suppressed beyond the correlation length ξ of the Debye plasma, causing pinch points to broaden at wave vectors $k \lesssim \xi^{-1}$: along directions of hitherto vanishing correlation functions, Lorentzian contributions of the form (1.28) appear. Measuring the broadening of pinch points therefore a sensitive probe of the presence of spinons [65–67].

1.3.2 Dipolar spin ice

The discussion of nearest-neighbour spin ice above captures the most important features of $Dy_2Ti_2O_7$ and $Ho_2Ti_2O_7$, except for the substantial m_J angular momentum components associated with the Ising variables σ. These give rise to effective magnetic moments $\mu_{\text{re}} \approx 10\mu_B$ on each rare-earth ion [45, 46, 68, 69] that point along the local $\langle 111 \rangle$ axes; due to the highly anisotropic g factor, magnetism normal to this direction is negligible. The magnetic fields of these large moments interact directly with one another, giving rise to long-range dipolar interactions comparable in strength to the short-range exchange interactions. These interactions are captured by the Hamiltonian [68]

$$H = \frac{J}{3} \sum_{\langle ij \rangle} \sigma_i \sigma_j + D\ell^3 \sum_{i<j} \sigma_i \sigma_j \left[\frac{\hat{e}_i \cdot \hat{e}_j}{r_{ij}^3} - \frac{3(\hat{e}_i \cdot \vec{r}_{ij})(\hat{e}_j \cdot \vec{r}_{ij})}{r_{ij}^5} \right], \tag{1.29}$$

where J is the strength of nearest-neighbour interactions, scaled by $1/3$ to account for the angles between local $\langle 111 \rangle$ axes; $D = \mu_0 \mu_{\text{re}}^2 / (4\pi\ell^3)$ is the energy scale of dipolar interactions,[†] normalised by the nearest-neighbour distance $\ell = a_0/\sqrt{8}$ of pyrochlore sites; $\vec{r}_{ij}$ is the vector pointing from spin i to spin j; $\hat{e}_{i,j}$ are unit vectors along the local $\langle 111 \rangle$ directions on the same spins such that $\sigma\hat{e}$ is parallel to their magnetic moments; the expression in square brackets is the space dependence of dipolar interactions [cf. (1.25)]. In $Dy_2Ti_2O_7$ and $Ho_2Ti_2O_7$, $\mu_{\text{re}} \approx 10\mu_B$ and $a_0 \approx 10$ Å, giving $D \approx 1.4$ K, similar to the exchange couplings $J = -3.72$ K in $Dy_2Ti_2O_7$ [68] and $J = -1.56$ K in $Ho_2Ti_2O_7$ [61]. Notably, the latter are antiferromagnetic and would lead to AIAO antiferromagnetic ordering on their own, requiring dipolar interactions to stabilise the spin-ice phase.

The most remarkable feature of this *dipolar spin-ice* model is that spinons become sources of the *physical* magnetic field, that is, *emergent magnetic monopoles* [62]. This is so because each vector of the lattice vector field introduced for the six-vertex model is now associated with a physical magnetic moment, which means that the coarse-grained

[†]The Ising variables σ are usually taken to be ± 1 rather than $\pm 1/2$ in the spin-ice literature.

$\vec{B}$ field is proportional to the local magnetisation of the material, the sources of which indeed double as sources of magnetic field.

A more microscopic and quantitative argument can be obtained from the *dumbbell model* [62]. The construction starts by replacing the magnetic moment of each atom with a pair of magnetic monopoles of magnitude $\pm q_0/2$, sitting on the centres of the two tetrahedra the spin belongs to (these are $a_\mathrm{d} = \sqrt{3}a_0/4$ apart), such that their net dipole moment, $q_0 a_\mathrm{d}/2$, equals the rare-earth magnetic moment [Fig. 1.7(c,d)]. The two configurations only differ in quadrupole and higher moments, the magnetic fields of which decay as $1/R^4$ or faster, which makes them negligible beyond nearest-neighbour spins.[†] Now, in the centres of 2120 tetrahedra, the monopoles from the four "dumbbells" cancel out [Fig. 1.7(c)]; on spinons (i.e., 3110 and 1130 tetrahedra), however, they sum to $\pm q_0$ [Fig. 1.7(d)]. That is, the magnetic field patterns of dipolar spin ice beyond atomic length scales are described accurately by magnetic monopoles of charge q_0, with Coulombic $1/R^2$ magnetic fields, associated with spinon sites. Accordingly, they exhibit an *energetic* Coulomb interaction, mediated by the same magnetic fields:

$$V(r_{ij}) = \begin{cases} \dfrac{\mu_0}{4\pi}\dfrac{q_i q_j}{r_{ij}} & i \neq j \\ V_0 q_i^2/2 & i = j; \end{cases} \qquad V_0 q_0^2 = \frac{4J}{3} + \frac{16}{3}\left[1 + \sqrt{\frac{2}{3}}\right] D; \tag{1.30}$$

the spinon chemical potential V_0 is chosen so as to recover the exact interaction between nearest-neighbour spins [62].

In $Dy_2Ti_2O_7$ and $Ho_2Ti_2O_7$, the magnitude of the emergent monopoles, $q_0 = 2\mu_\mathrm{re}/a_\mathrm{d}$, is about 8000 times smaller [62] than the vacuum monopole strength predicted by Dirac quantisation [74], $h/(\mu_0 e)$. This means that each pair of positive and negative monopoles are connected by a "Dirac string" [Fig. 1.7(e)] made up of the dumbbells flipped while separating them, which carries flux $\mu_0 q_0$ to compensate for the magnetic field emanating from the monopoles. In principle, these Dirac strings are observable; the massively degenerate 2120 ground-state manifold, however, makes them energetically irrelevant and their choice ambiguous.

At temperatures well below the spinon chemical potential $V_0 q_0^2$, monopole density is negligible, restricting spin configurations to the 2120 manifold: the latter is no longer exactly degenerate due to quadrupolar magnetic interactions, but this is irrelevant at all but the lowest temperatures [75]. As a result, dipolar spin ice shows the same long-range dipolar correlations (1.25, 1.26) and pinch points (1.27) as the nearest-neighbour model. Spinons introduced at higher temperatures reduce the contrast of these pinch points

[†] At the lowest temperatures, these effects would order dipolar spin ice in a particular 2120 state [70]; in practice, however, its dynamics becomes glassy before this transition could occur [71–73].

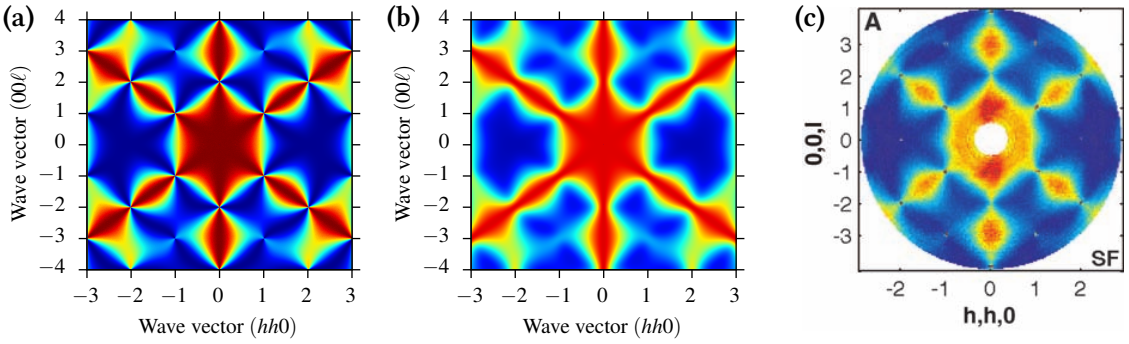

Figure 1.8. Pinch points in classical spin ice.
(a,b) Monte-Carlo simulations of nearest-neighbour spin ice (1.24) at $T = 0$ **(a)** and at $T = J/3$ **(b)**. The zero-temperature 2120 rule results in long-range correlations, manifest in sharp pinch points; at finite temperature, a thermal plasma of spinons introduces a finite correlation length, broadening the pinch points.
(c) Spin-flip neutron-scattering intensity in $Ho_2Ti_2O_7$. The pinch points are well visible; by measuring their blurring accurately, the density of spinons can be inferred. Figure (c) taken from Ref. 66.

too; however, the singularity is not blurred completely, only reduced in intensity [76]. To understand why this is so, consider a high-temperature expansion of the Boltzmann factors $e^{-\beta H[\sigma]}$: to lowest order, the correlations between spins are $\langle \sigma_i \sigma_j \rangle = -H_{ij}/T$, where H_{ij} is the coefficient of $\sigma_i \sigma_j$ in the Hamiltonian [75]. Since the dipolar spin-ice Hamiltonian (1.29) contains long-range interactions, spin correlations even at the highest temperatures retain a dipolar component that gives rise to identical-looking pinch points, merely suppressed in intensity as $1/T$.

1.3.3 Experimental approaches

The best-known signatures of spin-ice physics are pinch points due to long-range dipolar spin correlations. These can readily be measured using neutron-scattering experiments, on account of the coupling between the magnetic moments of the rare-earth ions and the neutrons. Such experiments are usually performed in the geometry first used in Ref. 61: the neutrons are polarised along the $[1\bar{1}0]$ direction and the momentum transfer is normal to the same, so it can be written as $\vec{q} = (hhk)$. In this geometry, the spin-flip channel [which probes the component of the tensor correlator (1.26) normal to both $\vec{q}$ and the polarisation axis] exhibits pinch points at several reciprocal lattice vectors [61] (Fig. 1.8). At finite temperatures, these pinch points broaden due to the presence of spinons that break the 2120 rule that underpin long-range correlations in the ground state [63, 66, 76]. Applying Debye–Hückel theory to the sparse, Coulomb in-

teracting plasma of spinons shows that this broadening is Lorentzian with a coherence length $\xi \propto \sqrt{\rho}$ [63]. Such blurring has been observed experimentally [Fig. 1.8(c)], and the inferred monopole density was found to have the expected Arrhenius temperature dependence [66].

Beyond correlation functions, the creation of gapped quasiparticle excitations is associated with a broad Schottky peak in the heat capacity: the integrated entropy change is consistent with the difference between the Pauling estimate (1.2) and the entropy of an infinite-temperature paramagnet, $Nk_{\mathrm{B}} \log 2$ [13, 59]. Furthermore, magnetic fields H along a $\langle 111 \rangle$ direction act as a staggered chemical potential for monopoles, giving rise to a liquid–gas type phase transition in (H, T) space [62, 77]. The recombination dynamics of spinons is also affected by their magnetic Coulomb interactions: the resulting changes in magnetic field fluctuations can be measured using muon spin rotation [78], allowing the monopole charge q_0 to be measured directly.

It has recently been proposed to observe the monopolar magnetic fields of spinons directly, by measuring the magnetic flux across a superconducting quantum interference device (SQUID) wrapped around a sample of spin ice [79–81]. A single monopole traversing the SQUID coil changes the net flux across the same by $\mu_0 q_0$, which can be measured using modern devices; similar methods have been used to search for Dirac-quantised magnetic monopoles in vacuum [82]. In spin ice, however, the density and recombination frequency of monopoles is too high to resolve them individually in this geometry: nevertheless, features in the noise spectrum consistent with both monopoles and the topological constraints that arise from the Dirac strings connecting them can be observed, providing a direct measurement of monopole dynamics [80]. Other geometries and experimental techniques, potentially able to resolve individual monopoles, have also been proposed [79].

2

Seeing beyond the light: Semiclassical simulations of quantum spin ice

Understanding the nature and behaviour of excitations is of great importance for defining, characterising, and classifying quantum spin liquids and topological phases in general. It has also proven crucial for experimental detection and characterisation of candidate materials [83]. Current theoretical and numerical techniques, however, have limited capabilities, especially when it comes to studying gapped excitations. Here, we propose a semiclassical numerical method to study systems whose spin-liquid behaviour is underpinned by perturbative ring-exchange Hamiltonians. Our method can readily access both thermodynamic and spectral properties. We focus in particular on quantum spin ice and its photon and vison excitations. After benchmarking the method against existing results on photons, we use it to characterise visons and their thermodynamic behaviour, which remained hitherto largely unexplored. We find that visons, in contrast to spinons in classical spin ice, form a weak electrolyte: vison pairs are the dominant population at low temperatures. This is reflected in the behaviour of thermodynamic quantities, such as pinch point motifs in the relevant correlators. Visons also hybridise strongly with the photon background, affecting how the former may show up in inelastic response measurements. Our results demonstrate that the method, and generalisations thereof, can substantially help our understanding of quasiparticles and their interplay in quantum spin ice and other quantum spin liquids, quantum dimer models, and lattice gauge theories in general.

2.1 Quantum spin ice (QSI)

> *All things pass and naught abides; you cannot step twice into the same stream.*
>
> — Heraclitus, as quoted in Plato's *Cratylus*

The classical spin-ice Hamiltonians (1.24, 1.29) are Ising models: all spin operators that make them up commute, and so they exhibit no quantum dynamics. Coupling between other spin components is, however, possible: the most general nearest-neighbour exchange Hamiltonian allowed by the symmetry of the pyrochlore lattice is of the form

$$H = \sum_{\langle ij \rangle} \Big\{ J_{zz}\sigma_i^z\sigma_j^z - J_\pm(\sigma_i^+\sigma_j^- + \sigma_i^-\sigma_j^+) + J_{\pm\pm}(\gamma_{ij}\sigma_i^+\sigma_j^+ + \gamma_{ij}^*\sigma_i^-\sigma_j^-) + J_{z\pm}\left[\sigma_i^z(\zeta_{ij}\sigma_j^+ + \zeta_{ij}^*\sigma_j^-) + i \leftrightarrow j\right] \Big\}, \tag{2.1}$$

where σ^z is the component of each spin pointing along its local $\langle 111 \rangle$ direction, the J are real parameters, and γ_{ij} and ζ_{ij} are unimodular complex numbers that depend on which FCC sublattices of the pyrochlore lattice the spins i and j belong to [50, 84]; nearest-neighbour classical spin ice is described by the first term only. The other three terms become relevant in pyrochlores where the overall spin of the rare-earth ion is relatively small, and the g-tensor of the low-energy effective doublet is not dominated strongly by the Ising component: Pr^{3+}, Nd^{3+}, Er^{3+}, and Yb^{3+} moments typically fit this description [50]. Both factors favour relatively small magnetic moments in the ground-state doublet, which makes long-range dipolar interactions and further-neighbour exchange interactions negligible; that is, the nearest-neighbour Hamiltonian (2.1) is expected to give a good account of all pyrochlores with significant quantum dynamics.

The only systematic approach to study the entire phase diagram of (2.1) to date is *gauge mean-field theory (GMFT)* [85–87], which is based on the understanding that the elementary excitations of classical spin ice are fractionalised spinons. In the presence of non-Ising terms, these spinons become dynamical, as $\sigma^\pm$ spin operators create a pair of oppositely charged spinons on two neighbouring tetrahedra, or move them from one to the other. GMFT separates these operators into a residual spin operator and explicit spinon-creation and annihilation operators that live on the tetrahedra, giving rise to an interacting spin–spinon Hamiltonian: using a self-consistent mean-field approximation, this Hamiltonian can be separated into a tight-binding spinon Hamiltonian coupled to static spins. Now, approximate ground states can be obtained by minising the energy of such a mean-field state. The resulting phase diagrams are shown in Fig. 2.1: they exhibit a variety of spin-liquid, spin-nematic, and unconventional ordered phases. Since

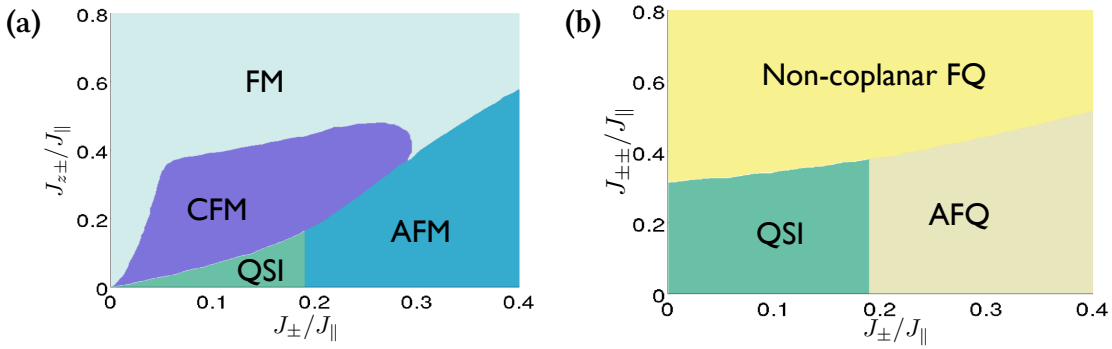

Figure 2.1. Phase diagram of the pyrochlore spin-1/2 Hamiltonian (2.1) obtained from gauge mean-field theory [50, 85–87]. The quantum-spin-ice (QSI) phase appears for small values of $J_\pm$ and $J_{\pm\pm}$, but is almost immediately destroyed by $J_{z\pm}$ interactions. It is surrounded by an exotic Coulomb ferromagnet (CFM) and ordered dipolar (FM, AFM) and quadrupolar (FQ, AFQ) phases. Figure taken from Ref. 50.

these phases are all derived from classical spin ice, and their signatures are often difficult to distinguish experimentally [88], it is common to refer to all phases in these phase diagrams without conventional magnetic order as quantum spin ice.

Out of these phases, I shall focus on a particular quantum spin liquid (green in Fig. 2.1), brought about by small transverse interactions on top of the dominant Ising terms. This phase, called *quantum spin ice (QSI)* in the stricter sense, is stable in the presence of finite $J_\pm$ and $J_{\pm\pm}$ interactions, but is destroyed by small $J_{z\pm}$ couplings. For simplicity, I shall focus on $J_\pm$ couplings only, resulting in the XXZ Hamiltonian [15]

$$H = \sum_{\langle ij\rangle}\left[J_z\sigma_i^z\sigma_j^z - J_\perp(\sigma_i^x\sigma_j^x + \sigma_i^y\sigma_j^y)\right] = \sum_{\langle ij\rangle}\left[J_z\sigma_i^z\sigma_j^z - \frac{J_\perp}{2}(\sigma_i^+\sigma_j^- + \sigma_i^-\sigma_j^+)\right]. \tag{2.2}$$

If $J_z > 0$ and $|J_\perp| \ll J_z$, transverse terms can be included as a perturbation acting within the ground-state manifold of the CSI defined by the Ising interactions. However, a single pair of $\sigma_i^+\sigma_j^-$ operators introduces a pair of spinons at next-nearest-neighbour tetrahedra: therefore, they can only act at higher orders in perturbation theory, by flipping spins around a closed loop (cf. §1.1.2). The shortest nontrivial loops are the hexagonal plaquettes shown in Fig. 2.2 that can be flipped in third order of perturbation theory; at small $J_\perp$, therefore, Eq. (2.2) is approximated well by the ring-exchange Hamiltonian

$$H_{\text{ring}} = -\frac{g}{2}\sum_{\hexagon}\left(\sigma_1^+\sigma_2^-\sigma_3^+\sigma_4^-\sigma_5^+\sigma_6^- + \text{H.c.}\right), \tag{2.3}$$

where $g = 3J_\perp^3/J_z^2$ is the perturbative amplitude of flipping the loop [15].

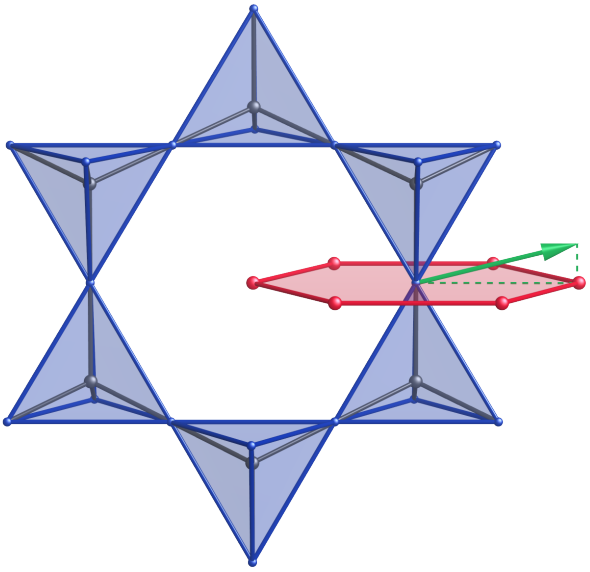

Figure 2.2. The pyrochlore lattice consists of the bond midpoints of a diamond lattice (black); due to the coordination of the latter, pyrochlore sites form corner-sharing tetrahedra (blue). The shortest closed loops of the pyrochlore lattice are regular hexagonal plaquettes whose centres (red dots) themselves form a dual pyrochlore lattice; the plaquettes of this lattice (red hexagon) are in turn centred on the original sites. In spin-ice models, Ising (CSI) or Heisenberg (QSI) spins live on each pyrochlore site. In our semiclassical simulation, the anisotropy of the Hamiltonian (2.3) results in mostly easy-plane spins for which σ^z (the component along the $\langle 111 \rangle$ direction of the corresponding diamond-lattice link) is a small fluctuation. Figure taken from Ref. 2 and created using VESTA 3 [43].

2.1.1 Ring exchange to lattice gauge theory

By construction, the ring-exchange Hamiltonian (2.3) conserves the spinon charge $q = \sum_{i \in t} \sigma_i^z$ of each tetrahedron t. As discussed in §1.1.2, this gives rise to U(1) gauge invariance under the local transformations $e^{i\chi q}$: indeed, nearest-neighbour classical spin ice is equivalent to the six-vertex model, and the ring-exchange term is identical to (1.7) apart from plaquette size. Accordingly, (2.3) can be rewritten exactly as a compact U(1) lattice gauge theory using the soft-spin representation $\sigma^z \to n$, $\sigma^\pm \to e^{\pm i\phi}$ [15], where n and ϕ are canonically conjugate operators with half-integer and angular values, respectively. Since the pyrochlore lattice can be regarded as the set of links of the bipartite diamond lattice formed by the tetrahedra, lattice electric fields $e_{rr'}$ and vector potentials $a_{rr'}$ can be defined on these links as

$$e_{rr'} = \begin{cases} +n & \text{if } r \in A, r' \in B \\ -n & \text{if } r \in B, r' \in A, \end{cases} \qquad a_{rr'} = \begin{cases} +\phi & \text{if } r \in A, r' \in B \\ -\phi & \text{if } r \in B, r' \in A, \end{cases} \tag{2.4}$$

where r and r' are diamond-lattice sites [i.e., centres of tetrahedra, cf. Fig. 1.6(b)], and A and B are the sublattices of the diamond lattice (that is, "up" and "down" tetrahedra). In fact, as $n \sim \sigma^z$ is the spin projection along the corresponding $\langle 111 \rangle$ axis, $e_{rr'}$ is the component of $\vec{\sigma}$ along the direction $\vec{r}' - \vec{r}$ (cf. Fig. 1.1). Now, following the steps of §1.1.2, the Hamiltonian (2.3) is equivalent to

$$H = \frac{U}{2} \sum_i e_i^2 - g \sum_{\hexagon} \cos \operatorname{curl} a, \tag{2.5}$$

where $U \to \infty$ to enforce the constraint $n = \pm 1/2$.

As mentioned in §1.1.2, this pathological limit can show deconfined physics analogous to that of continuum electrodynamics only by virtue of the frustrated half-integer electric fields e. To get a better handle on the physics, it is instructive to replace $\sigma^z = 1/2$ spins with dimers that join the centres of the two tetrahedra [15]: this maps σ^z configurations on the pyrochlore lattice to dimer coverings of the diamond lattice of tetrahedra, where each site belongs to two dimers; the ring-exchange Hamiltonian (2.3) takes the symbolic form

$$H_{\text{ring}} = -g \sum_{\hexagon} \big(|\hexagon\rangle\langle\hexagon| + |\hexagon\rangle\langle\hexagon| \big). \tag{2.6}$$

This dimer representation motivates introducing Rokhsar–Kivelson (RK) type potential terms (§3.2) that count the number of flippable plaquettes, leading to [19]

$$H = \sum_{\hexagon} \Big[-g\big(|\hexagon\rangle\langle\hexagon| + |\hexagon\rangle\langle\hexagon| \big) + \mu\big(|\hexagon\rangle\langle\hexagon| + |\hexagon\rangle\langle\hexagon| \big) \Big]. \tag{2.7}$$

At the RK point $g = \mu$, the ground state of (2.7) is an equal-amplitude superposition of all σ^z basis states that obey the two-in-two-out rule [89]: this is manifestly a U(1) quantum spin liquid (§3.2.1) and belongs to the same phase as any spin liquid that may form in the QSI case, $\mu = 0$. Indeed, this phase is expected to be stable for a finite range of $\mu \lesssim g$ (§3.2.1). Whether this includes $\mu = 0$ cannot be determined from field-theoretic arguments alone; however, efficient quantum-Monte-Carlo (QMC) methods can be used to obtain the ground state of (2.7) numerically, as it is not affected by the sign problem (§4.1.3).[†] These studies [19] show that the U(1) liquid phase extends from the RK point down to $\mu/g = -0.50(3)$, which includes the QSI Hamiltonian (2.6). Therefore, in the perturbative limit $|J_\perp| \ll J_z$, the ground state of the XXZ model (2.2) is a quantum spin liquid described by the deconfined phase of the compact U(1) lattice gauge theory (2.5).

The stability of this phase against finite $J_\perp$ can again be studied numerically. For

[†] This is manifest for $g > 0$; negative g introduces a $b = \pi$ magnetic flux on all plaquettes, which can be removed by an appropriate sign rule [15].

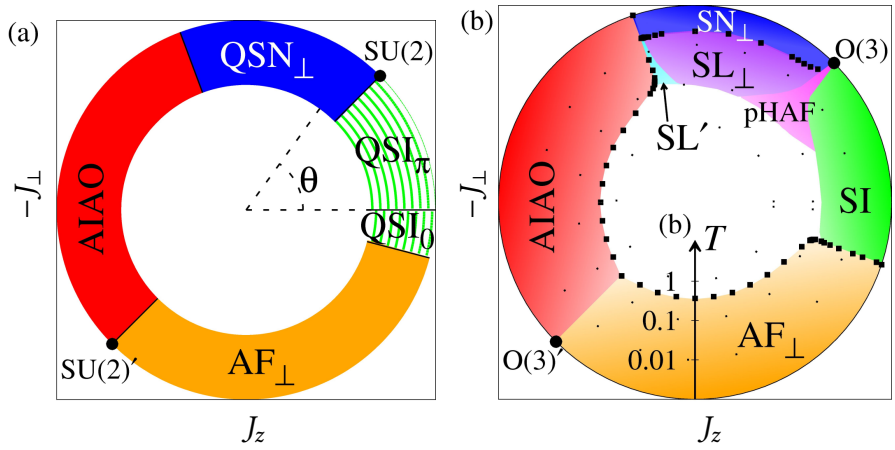

Figure 2.3. Phase diagram of the pyrochlore XXZ model (2.2).
(a) Zero-temperature phase diagram obtained from cluster-variational calculations. The QSI phase extends from $J_\perp \approx 0.104 J_z$ [90, 91] to the Heisenberg point $J_\perp = -J_z$; it is surrounded by XY antiferromagnetic ($\mathrm{AF}_\perp$) and spin nematic ($\mathrm{QSN}_\perp$) [88] phases. Ferromagnetic Ising interactions stabilise an all-in-all-out (AIAO) ordered phase.
(b) Thermal phase diagram of the XXZ model in semiclassical approximation. The QSI phase is replaced by classical spin ice (SI) between $-J_z < J_\perp < J_z/3$ at low temperatures. The critical end point $J_\perp = -J_z$ of the QSI phase extends into a pseudo-Heisenberg spin-liquid phase (pHAF) described by three U(1) gauge theories [92, 93]. A similar quantum critical phase forms at the other endpoint of the spin nematic phase (SL′). The other phases largely correspond to the quantum phase diagram.
Figure taken from Ref. 94.

positive $J_\perp$, (2.2) has no sign problem, allowing it to be studied using QMC techniques (cf. §4.1.3): these show a QSI phase extending up to $J_\perp = 0.104 J_z$, where it gives way to XY antiferromagnetic order [90, 91]. For $J_\perp < 0$, variational techniques [94] suggest that the QSI phase terminates at the antiferromagnetic Heisenberg point $J_\perp = -J_z$, which appears to be a critical point between QSI and a spin-nematic phase (Fig. 2.3).

2.1.2 Excitations of the U(1) gauge theory

The phenomenology of the deconfined phase of the U(1) gauge theory (2.5) is broadly similar to that of standard quantum electrodynamics. Small transverse fluctuations of the emergent electric and magnetic fields give rise to gapless photons; the spinons of classical spin ice become gauge charges by construction, with long-range Coulomb interactions mediated by the same photons. The compactness of the lattice gauge theory, however, provides for a straightforward construction of magnetic monopoles that obey Dirac quantisation [74] vis-à-vis spinons: we refer to these as *visons*.

Photons

In the deconfined phase of the U(1) gauge theory (2.5), the compactness of the vector potential a (or, equivalently, the discreteness of the electric field e) is irrelevant. Accordingly, the cosine term of (2.5) can be expanded to quadratic order around a given minimum, yielding the Hamiltonian [15, 17, 95]

$$H = \frac{U}{2} \sum_i e_i^2 + \frac{g}{2} \sum_{\hexagon} b^2; \tag{2.8}$$

where the magnetic field b is the lattice curl of a (Appendix A.1). It is important to point out that this expansion removes the frustration of the original compact theory: to remain in the deconfined phase, U has to be renormalised to a finite value on the order of g [95]. Now, e and b can be coarse-grained into continuum electric and magnetic fields $\vec{E}$ and $\vec{B}$, turning (2.8) into the standard Maxwell Hamiltonian (1.11). It follows that the excitation spectrum of QSI is gapless, with linearly dispersing photon modes: the speed of light is $c = \sqrt{\varepsilon\mu} \sim g a_0$, where a_0 is the pyrochlore lattice parameter.

Spinons

The only excitations of the nearest-neighbour classical spin-ice model (1.24) are spinons, which live on tetrahedra that violate the two-in-two-out rule. The construction of the ring-exchange Hamiltonian and the U(1) gauge theory promote the same rule into a gauge symmetry, turning spinons into gauge charges (§1.1.2). Indeed, the spinon charge (that is, the number of spins pointing out of the tetrahedron minus those pointing in) becomes the lattice divergence of the electric field e, as expected (Appendix A.1). As electric charges in the gauge theory (2.5), spinons acquire a $1/R$ Coulomb interaction mediated by the photons discussed above. The strength of this interaction is on the order of g between nearby spinons; it arises purely from the ring-exchange dynamics, without any explicit dipolar or Coulomb interaction in the Hamiltonian.

By construction, spinons are static under the ring-exchange Hamiltonian (2.3). Under the XXZ Hamiltonian (2.2), the transverse terms $\sigma^+\sigma^-$ cause the spinons to move, resulting in a spinon dispersion of width $\sim J_\perp$ around their energy under the classical spin-ice Hamiltonian, $J_z/2$.

Visons

An important feature of the gauge theory (2.5) is lost in taking the quadratic approximation (2.8). It is clear from (2.4, 2.5) that the vector potential a and the magnetic field

b are only defined up to 2π. As a result, the lattice analogue of $\operatorname{div}\operatorname{curl} a = 0$ only holds modulo 2π: a tetrahedron of the dual pyrochlore lattice can be a source of the emergent magnetic field, quantised in units of 2π.

The physical significance of such a source is clearest far away from it, where the field b is expected to be small and thus (2.8) holds. By Gauss' theorem, the flux across any closed surface of plaquettes surrounding the source will be 2π.[†] That is, in the electromagnetic formalism of (2.8), these defects can act as quantised physical sources of the no longer multi-valued magnetic field. The quantisation condition is equivalent to the Dirac quantisation of magnetic monopoles in free space [74].

These emergent "magnetic monopoles" are usually called *visons*: unlike $\mathbb{Z}_2$ visons, they are bosons and have trivial mutual statistics with the spinons. Just like spinons, they are gapped [by an $O(g)$ gap] and interact through an energetic Coulomb interaction mediated by photons. The vison gap is comparable to the bandwidth of photons and their interactions (all being controlled by g), so hybridisation between photons and visons away from the long-wavelength limit cannot be ruled out.

2.1.3 Naming conventions

> *I have always been obsessed with naming things. If I could name them, I could know them. If I could name them, I could tame them. They could be my friends.*
>
> — Eve Ensler, *The Vagina Monologues*

There is quite some ambiguity around the names of emergent gauge fields in the literature, due to the different approaches one can take to arrive at the quantum spin-ice model (2.3). In dipolar classical spin ice (and hence in real QSI candidate materials), spinons are effective sources of the physical magnetic field (§1.3.2) [62]; furthermore, it is conventional to describe the divergence-free classical fields arising in both classical spin ice (§1.1.1) and dimer models (§3.1.2) in terms of a magnetic field. This motivates calling e a *magnetic field* and referring to spinons as *magnetic monopoles.* Due to the Dirac duality between spinons and visons, the latter should then be thought of as an *electric charge,* turning b into an *electric field.*[‡]

[†] One may require a unique value for each a_i, hence strictly enforcing $\operatorname{div}\operatorname{curl} a = 0$. In this case, the net flux must be zero: the effect of a 2π-flux can be reproduced by threading a compensating flux across a single chain of plaquettes. Such a chain is the lattice analogue of the *Dirac string* for magnetic monopoles in a vacuum [74] but, just like a Dirac string, it has no observable physical consequences.

[‡] In fact, the possibility of direct coupling between b and the physical electric field has been proposed in Ref. 96, based on symmetry considerations.

This, however, stands at odds with the gauge theory convention that gauge charges (in this case, spinons) are sources of the *electric* field, and vector potentials give rise to *magnetic* fields. In this convention, e and b are electric and magnetic fields, respectively, while spinons and visons are electric charges and Dirac magnetic monopoles, respectively.[†]

Here and in chapter 3, I shall follow the latter convention, also used by Ref. 15, since constructing a field theory from first principles is crucial to both pieces of work. To reduce ambiguity, I shall refer to spinons and visons by these names, instead of electric or magnetic charges. For a brief history of these naming conventions and a survey of usage in the literature, see Ref. 50.

2.1.4 Signatures of the QSI phase

The most conventional probe of cooperative paramagnetism in pyrochlores is neutron scattering, which allows us to measure spin–spin correlation functions in reciprocal space. As discussed in §1.3.1, the long-range correlations borne out of the two-in-two-out constraint for classical spin ice manifest themselves as pinch points near reciprocal lattice points [12, 61, 66]: these pinch points are reproduced by the RK Hamiltonian (2.7) at $\mu = g$, where all 2120 states contribute equally [19]. By contrast, in the QSI ring-exchange Hamiltonian, scattering intensity near reciprocal lattice vectors is suppressed, "hollowing out" the pinch points [19, 95] (Fig. 2.4). Indeed, at zero temperature, each photon mode contains energy $\hbar\omega(\vec{q})/2$, which equals the energy expectation value of the corresponding electric and magnetic fields:

$$\frac{\hbar\omega(\vec{q})}{2} = \frac{\varepsilon\langle|E(\vec{q})|^2\rangle}{2} + \frac{\langle|B(\vec{q})|^2\rangle}{2\mu}. \tag{2.9}$$

As a result, the static spin correlator $\langle\sigma^z(\vec{q})\sigma^z(-\vec{q})\rangle$, which is the fine-grained form of $\langle|E(\vec{q})|^2\rangle$, is proportional to the dispersion relation $\omega(\vec{q})$; since this dispersion is gapless at the Γ point, the correlator is expected to vanish, too [95].

One can use the same argument to estimate static correlation functions at finite temperature, too [95]. As temperature increases, the occupation of each photon mode approaches equipartition, thus recovering classical-spin-ice pinch points once the temperature reaches a few times g (Fig. 2.4). This is accompanied by a Schottky peak in heat capacity [91, 98], which increases the entropy of the system to the Pauling value (1.2)

[†]The existence of both electric and magnetic charges makes the theory perfectly symmetric under swapping the words "electric" and "magnetic:" the Hamiltonian (2.5) can be recast in terms of the field b and a "dual vector potential" g such that $e = \operatorname{curl} g$, turning visons into *bona-fide* gauge charges [15]. Doing so, however, is usually inconvenient and is only done in practice to gain better access to visons [97].

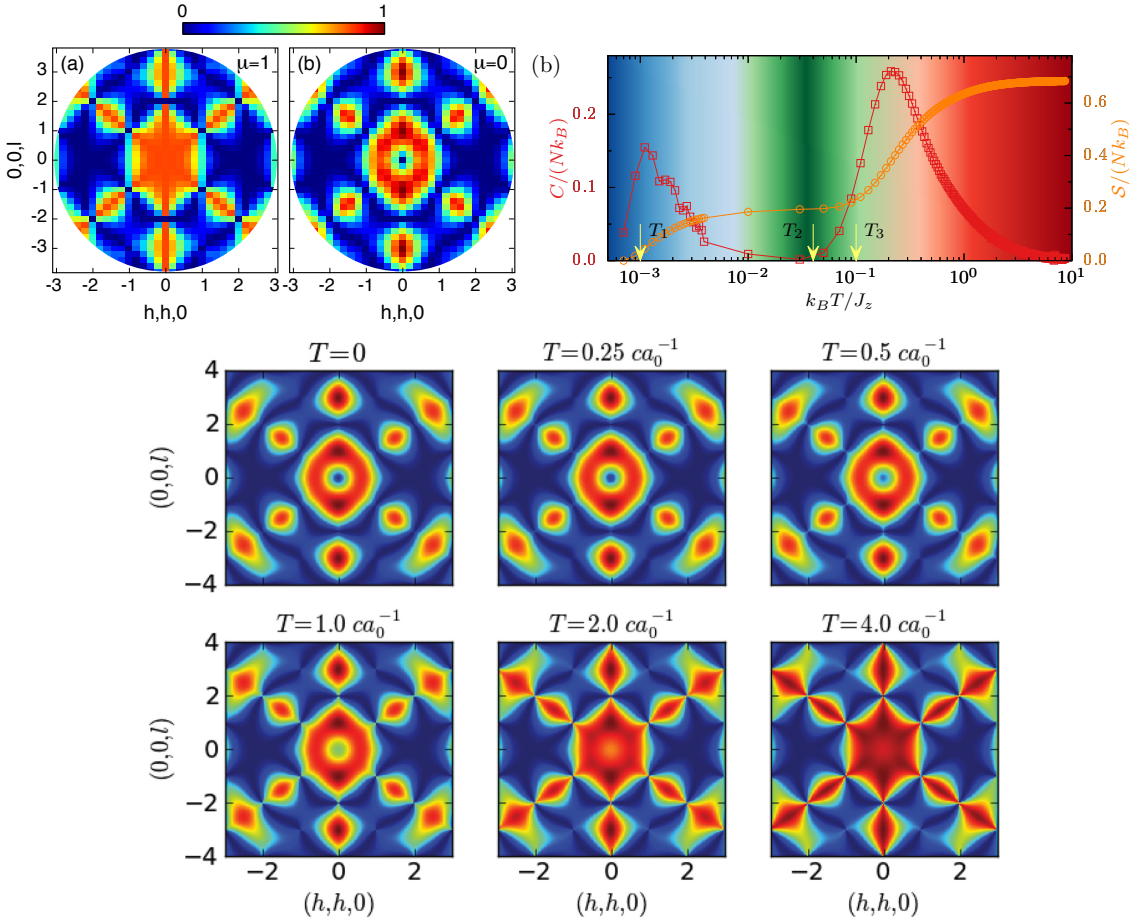

Figure 2.4. **Top left**: Zero-temperature spin-flip elastic neutron-scattering intensity of the RK Hamiltonian (2.7), in the setup of Ref. 66, obtained from quantum-Monte-Carlo (QMC) simulations. The pinch points prominent in classical spin ice [**(a)**: $\mu = g$] are suppressed in low-temperature QSI [**(b)**: $\mu = 0$]. Figure taken from Ref. 19.

Top right: Entropy (orange line) and heat capacity (red line) of the XXZ exchange Hamiltonian at $J_\perp = 0.092|J_z|$, obtained from QMC. The two broad peaks correspond to smooth crossovers from quantum spin ice (blue) to classical spin ice (green) to trivial paramagnet (red): the corresponding increases in entropy are consistent with the Pauling entropy formula (1.2). Figure taken from Ref. 98.

Bottom: Temperature dependence of the static structure factor (see above) in the ring-exchange Hamiltonian, obtained from field-theory calculations. The pinch points, suppressed at low temperatures, are gradually restored by a proliferation of photons at $O(g)$ temperatures. Figure taken from Ref. 95.

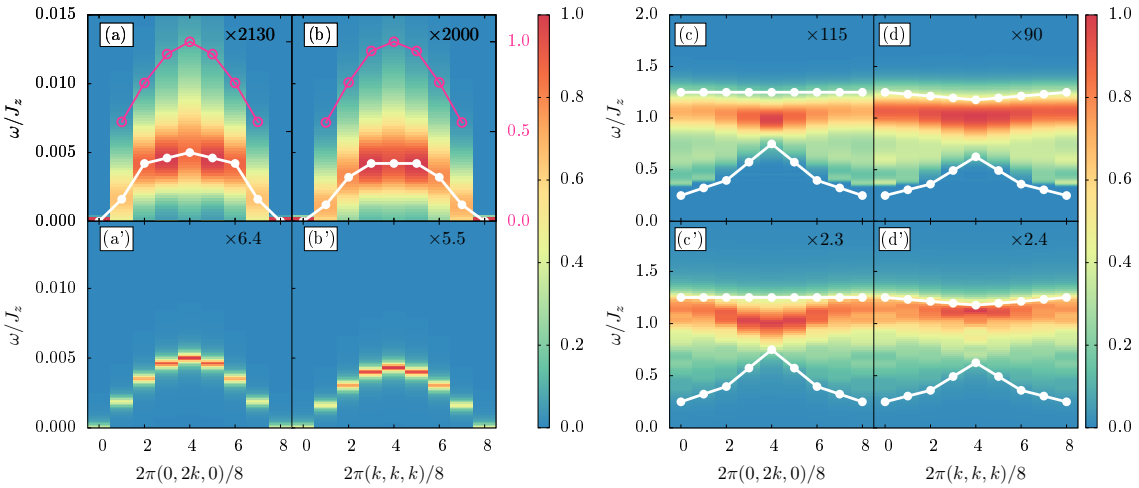

Figure 2.5. Quantum-Monte-Carlo (QMC) simulation of the dynamical structure factor in the XXZ Hamiltonian (2.2); $J_\perp = 0.092|J_z|$, $T = 0.001|J_z|$.
Left panel: $S^{zz}(\vec{q},\omega) = \langle\sigma^z(\vec{q},\omega)\sigma^z(-\vec{q},-\omega)\rangle$ from **(a, b)** QMC and **(a', b')** field theory. Photon modes are visible as a linearly dispersing branch of excitations, reaching zero energy at the Γ point.
Right panel: **(c, d)** $S^{+-}(\vec{q},\omega) = \langle\sigma^+(\vec{q},\omega)\sigma^-(-\vec{q},-\omega)\rangle$ from QMC. The diffuse continuum at $\omega \approx J_z$ is due to two-spinon excitations (no local operator can create single spinons). **(c', d')** Approximate structure factor obtained from a tight-binding model of spinons.
Figures taken from Ref. 98.

consistent with classical spin ice. For temperatures much greater than $g \sim J_\perp^3/J_z^2$, therefore, QSI behaviour is reduced to that of CSI. At $T \sim g$, photon and vison modes saturate simultaneously, resulting in a single Schottky peak in which features characteristic for each (e.g., Debye heat capacity for photons vs. thermal activation for visons) cannot be separated cleanly. A possible experimental strategy around this problem is focusing on thermal transport [99], where the Arrhenius scaling of visons becomes qualitatively distinct from the $1/T$ photon thermal conductivity.

Due to the nontrivial quantum dynamics of QSI, dynamical correlation functions offer, in principle, more detailed signatures of QSI behaviour. In particular, the fine-grained electric-field correlator $\langle\sigma^z(\vec{q},\omega)\sigma^z(-\vec{q},-\omega)\rangle$ is expected to show sharp, linearly dispersing photon modes at least at small wave vectors [cf. (2.9)], while $\langle\sigma^+(\vec{q},\omega)\sigma^-(-\vec{q},-\omega)\rangle$ would display a diffuse two-spinon continuum at $\omega \approx J_z$, broadened by $\sim J_\perp$ on account of spinon dynamics under the XXZ Hamiltonian, since $\sigma^\pm$ creates a pair of spinons.[†] These features were demonstrated using QMC simulations of (2.2) [98], see Fig. 2.5; the

[†]Note that single spinons cannot be observed using any measurement, as they are fractionalised excitations, always created or annihilated in pairs.

substantial broadening of photon modes at high wave vectors may be due to either photon–photon interactions, or photon–vison hybridisation. Picking up these features experimentally is, however, rather challenging: in practice, J_z is on the order of kelvins, while $J_\perp < J_z$ in the QSI phase; accordingly, $g \ll 1$ K, making the photon dispersion relation impossible to resolve in neutron-scattering experiments. Therefore, the (much less distinctive) broadening of the two-spinon continuum is typically used as the only dynamical probe of the QSI phase [47, 100].

2.1.5 Candidate materials

While no material has yet been confirmed as U(1) quantum spin ice unambiguously, several pyrochlores show signatures of quantum-spin-liquid behaviour. In recent literature [50], the most discussed candidates have been $Tb_2Ti_2O_7$ [101, 102], $Yb_2Ti_2O_7$ [84, 103, 104], and $Pr_2X_2O_7$ (X = Zr, Sn, Hf) [47, 99, 100, 105].

Neutron-scattering experiments on $Tb_2Ti_2O_7$ have shown pinch-point-like features, confirming substantial antiferromagnetic Ising coupling [101]. These pinch points are, however, elongated arm-like features rather than single points, suggesting deviations from the conventional classical-spin-ice picture. It has been demonstrated [102] that the classical thermodynamics of a Hamiltonian of the general form (2.1) can account for these "pinch-line" features.[†] Furthermore, the presence of any pinch points at low temperatures are at odds with their suppression in the QSI model (Fig. 2.4). In summary, it appears unlikely that $Tb_2Ti_2O_7$ be a QSI, if at all a QSL.

Neutron scattering in $Yb_2Ti_2O_7$ displays rods of intensity along $\langle 111 \rangle$ directions, confirming strong anisotropy which is, however, rather different from that in CSI [103]. Among all QSI candidates, $Yb_2Ti_2O_7$ has the most reliable measurements of the exchange parameters in (2.1) [84]: for these parameters, methods including classical and gauge mean-field theory [84–86], RPA, and classical spin dynamics [104] all predict a canted (with respect to the local $\langle 111 \rangle$ axes) ferromagnetic order, with a gapless *magnon* excitation spectrum. More recently, it was argued that the discrepancy between this ordered state and the experimental neutron-scattering pattern is due to competition with an easy-plane ordered state, made possible by strong magnon interactions [51, 106].

Currently, the most likely candidates for displaying the U(1) QSI phase are praseodymium-based pyrochlores. While $Pr_2Ti_2O_7$ is not stable in the pyrochlore structure, neutron-scattering measurements in $Pr_2Zr_2O_7$ [47] and especially $Pr_2Hf_2O_7$ [100] show a suppression of pinch points in the elastic scattering pattern as well as a gapped contin-

[†]Such a model can be physical, as crystal field effects may slow the spin dynamics down and hence render it classical, as is the case with CSI itself [73].

uum of excitations in the inelastic spectrum, interpreted as spinons (cf. Fig. 2.5). Furthermore, thermal-conductivity measurements in $Pr_2Zr_2O_7$ [99] show a combination of two Arrhenius relations with different activation energies, possibly corresponding to visons and spinons. These results are quite promising; however, they are limited by structural disorder of $Pr_2Zr_2O_7$ single crystals [100, 105], marking out $Pr_2Hf_2O_7$ as the most likely clean example of a U(1) quantum spin ice.

2.2 Semiclassical simulation of magnetic systems

There is always a well-known solution to every human problem – neat, plausible, and wrong.

— Henry Louis Mencken

Larmor precession is perhaps the earliest quantitative description of the dynamics of atomic and nuclear magnetism. It relies on the observation that such microscopic magnetic moments are tied to the angular momentum of the particle by the *gyromagnetic ratio* γ: $\vec{\mu} = \gamma\vec{J}$. In an external magnetic field $\vec{B}$, a torque is exerted on this moment, leading to the dynamical equation

$$\frac{\mathrm{d}\vec{J}}{\mathrm{d}t} = \vec{\tau} = \vec{\mu} \times \vec{B} = \gamma\vec{J} \times \vec{B}: \tag{2.10}$$

$\vec{J}$ precesses around the direction of the magnetic field with the *Larmor frequency* $\omega = \gamma B$. Using AC magnetic fields, resonances with this precession can be observed in ferromagnets [107] and nuclear spins [108], both of which give rise to important experimental techniques in condensed-matter physics and chemistry.

The Larmor precession formula (2.10) and its descendants (e.g., the Landau–Lifshitz–Gilbert equation [107]) are typically used to account for the net magnetic moment of a ferromagnetic bulk, rather than individual spins; nevertheless, (2.10) does hold for single spins in an external field. Furthermore, one can rewrite a generic Hamiltonian defined on Heisenberg spins,

$$H = -\frac{1}{2}\sum_{i\neq j} H_{ij}^{\alpha\beta}\sigma_i^\alpha\sigma_j^\beta, \tag{2.11}$$

in terms of effective *exchange fields* acting on each spin:

$$H \equiv -\frac{1}{2}\sum_i \vec{B}_i^{\mathrm{eff}} \cdot (\gamma\hbar\vec{\sigma}_i) \implies B_i^{\mathrm{eff},\alpha} = \frac{1}{\gamma\hbar}\sum_{j\neq i} H_{ij}^{\alpha\beta}\sigma_j^\beta. \tag{2.12}$$

To allow for the most generic bilinear interactions possible, the coupling between spins

i and j is taken to be a tensor H_{ij}, defined such that $H_{ij} = H_{ji}^{\mathrm{T}}$; summation over the Greek spin-component indices is implied. In the semiclassical approximation, we replace the effective field (2.12) with its expectation value and plug it into (2.10) to obtain the dynamical equations

$$\frac{\mathrm{d}(\hbar\vec{\sigma}_i)}{\mathrm{d}t} = \gamma\hbar\vec{\sigma}_i \times \vec{B}_i \implies \frac{\mathrm{d}\sigma_i^\alpha}{\mathrm{d}t} = \frac{1}{\hbar}\varepsilon^{\alpha\beta\gamma}\sigma_i^\beta \sum_{j\neq i} H_{ij}^{\gamma\delta}\sigma_j^\delta. \tag{2.13}$$

Now, treating the $\vec{\sigma}_i$ as precessing unit vectors, the dynamics of an arbitrary bilinear spin Hamiltonian can be simulated by evolving the ordinary differential equations (2.13). The computational effort (time or memory) needed for such a simulation is proportional to the number of spins: this makes the approach very attractive for accessing large system sizes in comparison with solving Schrödinger's equation exactly, which requires exponentially scaling resources.

The dynamical equations (2.13) can also be obtained from (2.11) without reference to Larmor precession, using Ehrenfest's theorem: the expectation value of each spin component σ_i^α evolves as

$$\frac{\mathrm{d}\langle\sigma_i^\alpha\rangle}{\mathrm{d}t} = \frac{1}{i\hbar}\left\langle\left[\sigma_i^\alpha, H\right]\right\rangle = \frac{1}{\hbar}\varepsilon^{\alpha\beta\gamma}\sum_{j\neq i} H_{ij}^{\gamma\delta}\left\langle\sigma_i^\beta\sigma_j^\delta\right\rangle; \tag{2.14}$$

Eq. (2.13) follows by replacing the expectation values $\langle\sigma_i^\beta\sigma_j^\delta\rangle$ with the products $\langle\sigma_i^\beta\rangle\langle\sigma_j^\delta\rangle$. For quantum spin systems, this step is a gross approximation, as it eliminates any quantum correlations or entanglement, both of which give rise to connected spin correlators $\langle\vec{\sigma}_i\vec{\sigma}_j\rangle - \langle\vec{\sigma}_i\rangle\langle\vec{\sigma}_j\rangle$. The accuracy of the approximation, however, improves as the spin quantum number S gets larger, since the connected correlators are $O(S)$ in the large-S limit, compared to the $O(S^2)$ spin products: accordingly, approaches using the dynamical equation (2.13) are referred to as *semiclassical* or *large-S dynamics*.

Quantum effects beyond the semiclassical description are strongest in spin-1/2 systems; as such, it would be reasonable for (2.13) to get their true quantum dynamics completely wrong. It is remarkable, therefore, that semiclassical approaches work well (at least qualitatively) in this limit, both in ordered and spin-liquid phases: see, for instance, Refs. [88, 92, 93, 104, 109–112] on pyrochlore and kagome antiferromagnets and Refs. 113–116 on Kitaev honeycomb systems. In ordered magnets, this is so because every spin develops a nonzero expectation value in the ground state, proportional to the order parameter, which is in turn an average over all spins in the system, and thus effectively classical: spin dynamics is dominated by Larmor precession due to the effective background field (2.12) of these classical spin components. This is the basis of *linear spin-*

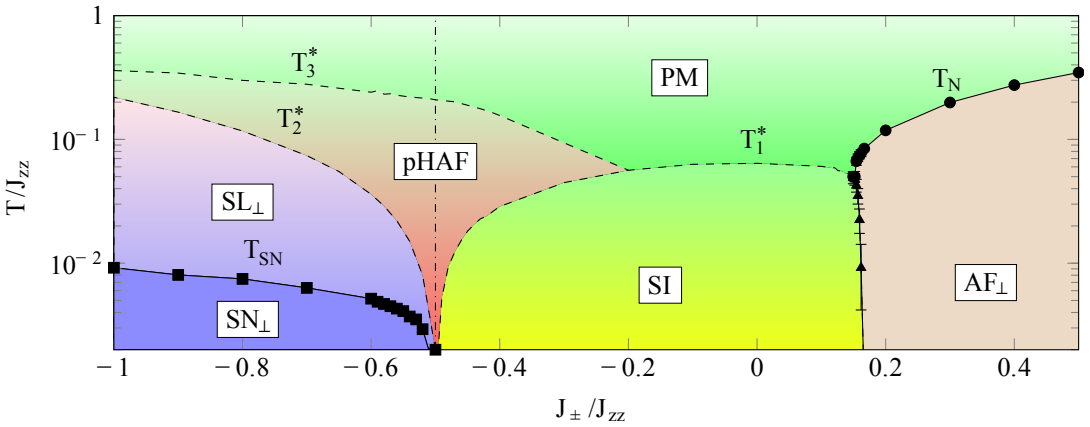

Figure 2.6. Classical phase diagram of the XXZ Hamiltonian (2.2). For $-|J_z| < J_\perp < |J_z|/3$, the model displays a *classical* spin-ice (SI) phase at low temperatures: ice-rule obeying states maximally polarised along the Ising axes are stable equilibrium states without any quantum dynamics between them. On the XY ferromagnetic side, the CSI phase transitions into an easy-plane ordered phase ($AF_\perp$); on the antiferromagnetic side, several new spin-liquid phases appear (see also Fig. 2.3). Figure taken from Ref. 88.

wave theory [117]: the full dynamical equation (2.14) can be linearised around the same ordered state, giving rise to normal modes known as *magnons;* their dispersion relation is captured exactly by semiclassical simulations. Likewise, spin liquids in Heisenberg models are often amenable to a bosonic parton-mean-field construction (§1.2): these can be extended straightforwardly to the large-S limit, which again reduces fluctuations around the mean field, retaining the spin-liquid physics in the semiclassical case. In higher orders of the $1/S$ expansion underlying semiclassics, however, spinon interactions can drive ordering transitions in $S = \infty$ spin-liquid phases [118]: applying semiclassics in these situations may lead to qualitatively wrong predictions. In general, the renormalisation and broadening of spinon dispersions and other interaction effects tend to be more significant than in linear spin-wave theory.

2.3 Semiclassics for ring-exchange dynamics

People get a lot of confusion because they keep trying to think of quantum mechanics as classical mechanics.
— Sidney Coleman

A crucial shortcoming of semiclassical simulations based on (2.13) is that they cannot access the ring-exchange dynamics of perturbative quantum spin liquids, such as QSI. The spin-liquid physics of these systems is borne out of an extensive number of degen-

erate ground states, coupled by dynamical terms that are dominated by ring exchange in the limit of small transverse coupling; this degeneracy, however, relies on small local Hilbert spaces. In the case of QSI, for instance, σ^z eigenstates with $\sum_i \sigma_i^z = 0$ (i.e., 2120 states) are degenerate under the six-vertex Hamiltonian (1.1) because $(\sigma_i^z)^2 = 1/4$ in both Ising states. This is no longer true if $S > 1/2$, lifting the degeneracy of the classical-spin-ice ground-state manifold, which disrupts the QSI phase. In particular, the lowest-energy basis states are CSI states with spins maximally polarised along the Ising axis [88, 119], as these minimise all terms in (1.1); accordingly, the QSI phase is replaced by static, *classical* spin-ice configurations in the semiclassical phase diagram of (2.2), see Fig. 2.6. Dynamics becomes restricted to small transverse fluctuations around these polarised states; ring-exchange processes only appear at orders $O(S)$ in pertubation theory, making them completely irrelevant.

By contrast, the ring-exchange Hamiltonian (2.3) remains in a spin-liquid phase even in the large-S limit [17]. This is possible because the Hamiltonian consists only of $\sigma^{\pm}$ operators, therefore, its classical ground state maximises the transverse components of spins. It follows that $\sigma^z = 0$ in equilibrium, which allows small fluctuations in both σ^z and the transverse angle ϕ: as these quantities are analogous to the lattice electric field and vector potential (2.4), normal modes of these fluctuations reproduce the photon excitations of QSI [17].

2.3.1 Dynamical equations

The time dependence of the expectation values $\langle \vec{\sigma}_i \rangle$ can be derived from Ehrenfest's theorem, as discussed in §2.2. Alternatively, we can start from a path-integral description, based on the action

$$\mathcal{S} = \hbar S \sum_i \int_0^T \mathrm{d}t \int_0^1 \mathrm{d}s\, \vec{n}_i \cdot (\partial_t \vec{n}_i) \times (\partial_s \vec{n}_i) - \int_0^T \mathrm{d}t\, H[\vec{n}(t)], \tag{2.15}$$

where the first term is the Berry phase for real-time evolution of spins [32],[†] and $\vec{n}_i$ is a unit vector that specifies the spin coherent state;[‡] the formally identical Hamiltonian $H[\vec{n}]$ is understood to act on unit vectors $\vec{n} \equiv \{\vec{n}_i\}$, rather than quantum spins. Vary-

[†] $\vec{n}(t,s)$ is smooth function in both variables such that $\vec{n}(t,1)$ is the trajectory $\vec{n}(t)$ and $\vec{n}(t,0)$ is a fixed "pole;" the integral gives the solid angle enclosed by the closed trajectory $\vec{n}(t)$, so it does not depend on the choice of this pole or the details of $\vec{n}(t,s)$ [32].

[‡] This is the eigenstate of the spin operator $\vec{n} \cdot \vec{\sigma}$ with the largest eigenvalue.

ing (2.15) yields the least-action trajectories

$$\hbar S\vec{n}_i \times (\partial_t \vec{n}_i) = \frac{\partial H[\vec{n}]}{\partial \vec{n}_i} \quad \Longrightarrow \quad \partial_t \vec{\sigma}_i = -\frac{1}{\hbar}\vec{\sigma}_i \times \frac{\partial H[\vec{\sigma}]}{\partial \vec{\sigma}_i}, \tag{2.16}$$

equivalent to Ehrenfest's theorem. The path-integral formalism, however, makes changing the parametrisation of the action, and hence the classical trajectories, straightforward, allowing us to rigorously relate our semiclassical simulation to other approaches. In particular, in the Villain spin representation in terms of σ^z and the transverse angle ϕ [120], the large-S action corresponding to (2.3) is manifestly equivalent to a U(1) lattice gauge theory [17]; this implies that the semiclassical trajectories (2.16) are equivalent to classical electromagnetism in the long-wavelength limit.[†]

For the ring-exchange Hamiltonian (2.3), the effective exchange field acting on each spin is given by

$$-\frac{\partial H[\vec{\sigma}]}{\partial \vec{\sigma}_i} \equiv \vec{h}_i^{\text{eff}} = (\operatorname{Re} h_i, \operatorname{Im} h_i, 0); \qquad h_i = g \sum \sigma^+_{i+1}\sigma^-_{i+2}\sigma^+_{i+3}\sigma^-_{i+4}\sigma^+_{i+5}, \tag{2.17}$$

where the summation is over the six plaquettes that i belongs to (Fig. 2.2) and $i{+}n$ denotes the nth spin counted from i on each plaquette (the direction around the plaquette is immaterial). Simulating the large-S dynamics of QSI now involves solving the differential equation (2.17); however, with no quasiparticles and infinitesimal zero point fluctuations, that solution would be trivial at $T = 0$, requiring simulations at finite temperatures.

2.3.2 Thermodynamics

In the path-integral representation, the massive entanglement that defines quantum spin liquids [7] appears as the interference of a continuum of equivalent trajectories, one related to the other by gauge symmetry [32]. The interference itself is a defining feature of spin-1/2 QSLs, since it accounts for differences between related Ising and quantum spin liquids [95]. For large S, however, quantum fluctuations become unimportant and interference effects are only apparent at the lowest temperatures. Therefore, a large-S QSL is generally indistinguishable from a classical spin liquid (CSL) characterised by a massive degeneracy of the least-action trajectories of the quantum path integral (2.15).

It is thus straightforward to obtain static correlation functions of the large-S QSL by Monte-Carlo sampling the CSL Boltzmann distribution given by $e^{-\beta H[\vec{n}]}$. Such sam-

[†]By contrast, the least-action trajectories of the large-S XXZ Hamiltonian describe fluctuations around the maximally polarised CSI ground states. Dynamics between the CSI states, equivalent to QSI physics, comes from instanton events *between* these trajectories, which cannot be captured using semiclassical simulations.

pling can be done more efficiently than quantum Monte Carlo and never suffers from a sign problem; furthermore, it naturally captures gapped excitations, which are hard to treat in analytic large-S calculations [17, 97]. Finally, using the same Monte-Carlo configurations as initial conditions for time evolution is a natural prescription for finding dynamic correlators of the classical spin liquid: in the large-S path integral language, this essentially samples least-action trajectories of a finite-temperature (e.g., Keldysh) path integral.

2.3.3 Implementation for QSI

To generate thermally distributed classical spin configurations, we use a Monte-Carlo algorithm that samples σ^z and the phase of σ^+ (for its magnitude is fixed by σ^z) independently – this mirrors the anisotropy of the Hamiltonian (2.3). Furthermore, we insist that there be no spinons in the system, that is, $\sum \sigma^z = 0$ for all tetrahedra of the pyrochlore lattice. This can be achieved by updating σ^z only in closed loops with alternating signs, similar to typical low-temperature simulations of CSI [70, 121]; as σ^z is now a continuous variable, loops can always be updated by small amounts. For convenience, we only perform updates around hexagons;[†] the proposed change is drawn from a Gaussian distribution whose variance is proportional to temperature (this ensures a good acceptance rate at all temperatures); acceptance is decided using the Metropolis method.

On the other hand, there are no conservation laws to be obeyed by the phase ϕ of σ^+, so it can be updated independently for each spin. In particular, since the magnitude of σ^+ is kept constant, ϕ follows the von Mises distribution

$$f(\phi_i) \propto e^{\beta \operatorname{Re}(\sigma_i^+ h_i^*)} = \exp\left[|\sigma_i^+|\,|h_i|\,\cos(\phi_i - \arg h_i)\right], \tag{2.18}$$

where the effective exchange field h_i is given by (2.17). This distribution can efficiently be sampled directly [122, 123], allowing ϕ_i to be sampled without rejection. Furthermore, to eliminate spurious correlations due to not sampling the U(1) gauge freedom, each Monte-Carlo step includes rotating ϕ for the four spins of each tetrahedron by an angle χ drawn from a uniform distribution between 0 and 2π.

[†]One can see that the hexagon updates lead to an ergodic Monte-Carlo protocol within each σ^z magnetisation sector as follows. The no-spinon constraint implies that σ^z is the sum of a pure lattice curl and a global magnetisation (Appendix A.2). Therefore, for any two valid configurations of σ^z with equal overall magnetisation, there is an "electric vector potential" G on the dual pyrochlore lattice such that $\operatorname{curl} G = \Delta\sigma^z$. Changing σ^z around each hexagon by this G on the corresponding dual pyrochlore site takes the system from one configuration to the other. In our work, we focus on the entropically dominant zero magnetisation sector only, as none of the properties we look at depend on changes in the macroscopic magnetisation of the system. If needed, winding loop updates (cf. §3.1.2) to sample different sectors could easily be introduced in the algorithm to remove this limitation.

Once initial spin configurations are sampled, they can be evolved in time using the dynamical equations (2.17). We integrated these numerically using the GNU Scientific Library implementation of the Prince–Dormand (8,9) ODE solver with automatic step-size control [124]. We obtain dynamical correlation functions of the semiclassical spin liquid by averaging over several initial configurations (§2.3.2).

It is important to note that, in the semiclassical picture, quantum correlations and uncertainties are irrelevant and therefore all quantities can be represented by pure numbers rather than quantum operators. That is, given a configuration of spins, the value of any observable can be determined unambiguously and straightforwardly. For intricate quantities, such as the magnetic field b, this is a substantial improvement over standard methods, e.g., quantum Monte Carlo. In particular, we follow the soft-spin prescription to identify $\sigma^{\pm}$ with $e^{\pm ia}$ [15] and take the vector potential a to be the complex argument of $\sigma^{+} = \sigma^{x} + i\sigma^{y}$. Now, the magnetic field b follows as

$$b = \operatorname{curl} a = \arg(\sigma_1^{+}\sigma_2^{-}\sigma_3^{+}\sigma_4^{-}\sigma_5^{+}\sigma_6^{-}); \tag{2.19}$$

the argument function is restricted to the interval $[-\pi, \pi)$ for a unique b with the smallest possible modulus, consistent with Ref. 15. Furthermore, as b is a lattice vector field, it has a sign ambiguity related to the orientation of the dual-diamond-lattice link representing the plaquette (see also Ref. 95 and Appendix A.1): we fix this ambiguity by orienting all links from "up" to "down" tetrahedra of the dual lattice, which is achieved by the proper choice of σ_1 on each plaquette.

2.4 Results

If it's stupid but it works, it isn't stupid.
— Murphy's laws of combat

In this section, I discuss the effectiveness of our semiclassical method applied to quantum spin ice. We first obtained the dispersion relation of the emergent photons from low-temperature dynamical simulations in §2.4.1: this compares excellently to theoretical results in the large-S limit [17], providing a useful benchmark on the technique. To further demonstrate the capabilities of the technique, we also studied static properties of QSI using Monte-Carlo sampling on its own. In particular, we measured the magnetostatic interaction between isolated visons, which can be inserted by hand in the simulations (§2.4.2). We find that visons, while deconfined, form a *weak electrolyte:* nearby pairs of positively and negatively charged visons have a lower energy than isolated ones, making their separation thermodynamically unfavourable. Finally, in §2.4.3,

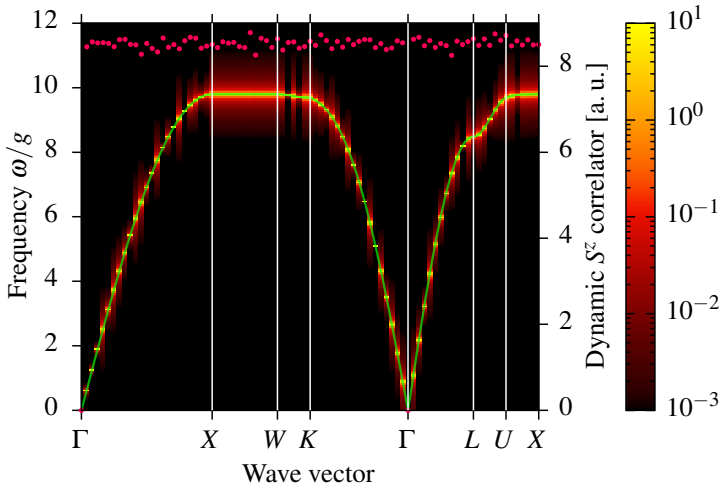

Figure 2.7. Dynamic structure factor $S^{zz}(\vec{q},\omega)$ along high-symmetry directions in the semiclassical QSI model at $T = 10^{-4}\,g$. Photons manifest as a sharp, gapless, linearly dispersing branch of classical normal modes. The frequency of these modes matches large-S analytic predictions (green line) excellently [17]. The integrated structure factor of the modes (red dots) is independent of $\vec{q}$, as expected on grounds of equipartition. Figure taken from Ref. 2. Perceptionally uniform colour map chosen following Ref. 125.

I discuss the interplay between visons and photon modes at high temperatures, as well as the thermodynamic consequences of the strong vison–vison interaction.

2.4.1 Photons

In order to obtain clean photon modes, it is important to minimise their interaction with one another and with other excitations. Spinons are excluded altogether by the loop-update Monte-Carlo algorithm; visons and photon–photon interactions can be eliminated by reducing the temperature. Importantly, as σ^z and ϕ are continuous variables, Monte-Carlo simulations can be run at arbitrarily small temperatures without increasing mixing times.

We generated 4096 stochastically independent Monte-Carlo configurations of $24 \times 24 \times 24$ cubic unit cells of the pyrochlore lattice (221 184 spins) at temperature $T = 10^{-4}\,g$ and calculated the time evolution of each for 2048 time steps of size $\delta t = (16g)^{-1}$. The results were Fourier transformed using the FFTW library [126] in time and space, separately for the four FCC sublattices of the pyrochlore lattice. Following Ref. [98], we evaluated the correlator

$$S^{zz}(\vec{q},\omega) = \sum_{\mu} \left\langle \sigma^z_\mu(\vec{q},\omega)\sigma^z_\mu(-\vec{q},-\omega)\right\rangle \tag{2.20}$$

along high-symmetry directions, where the sum runs over the sublattices; the results are plotted in Fig. 2.7. A single set of sharp normal modes appear in the data;[†] their frequencies match analytic results for the large-S photon dispersion perfectly [17] (green line in Fig. 2.7), confirming that the model simulated by our method is indeed equivalent to large-S QSI.

We note that the integrated structure factor $S^{zz}(\vec{q}) = \int \mathrm{d}\omega\, S^{zz}(\vec{q},\omega)$ is independent of $\vec{q}$, contrasting $S^{zz}(\vec{q}) \propto \omega(\vec{q})$ in spin-1/2 QSI [95]. This discrepancy is caused by the different realisation of photons in the two systems. In the quantum limit, zero-point fluctuations of the photon modes with energy $\hbar\omega(\vec{q})/2$ give rise to dynamic correlators even at zero temperature; in the large-S classical spin liquid, photons are classical normal modes with energy determined by equipartition, leading to $\vec{q}$-independent correlators.[‡]

2.4.2 Vison magnetostatics

U(1) visons in QSI are 2π-quantised sources of the emergent magnetic field b. Their existence and quantisation is due to the 2π-ambiguity of the transverse phases ϕ that are promoted to the vector potential a in the gauge-theoretic description. For the same reason, however, specifying vison numbers unambiguously is far from trivial. Typically, visons are understood through their far-field effects, where $b \sim 1/R^2$ is small and thus well-defined, and the total flux across a large, closed surface gives a unique vison charge [7]. In principle, one can define a U(1) vison charge operator

$$q = \operatorname{div} b/2\pi \tag{2.21}$$

for each dual diamond site by giving b on each plaquette a unique value (e.g., by restricting it to between $-\pi$ and π). However, since visons are dynamical, the system will not normally be in an eigenstate of q, which makes pinpointing them complicated. A great advantage of the semiclassical method is that observables like b and q are pure numbers rather than quantum operators: this means that the vison charge as defined above is always an unambiguous integer for all dual diamond sites.

The semiclassical simulation also allows us to introduce visons at will. We achieve this through a "fundamental step" that consists of rotating the transverse components of the six spins around a given plaquette by $\pi/4$ in alternating directions. Doing so changes

[†] The period of oscillations is not necessarily commensurate with the simulation time window. After a discrete Fourier transform, this leads to the slight broadening of the dispersion seen in Fig. 2.7.

[‡] In terms of the large-S expansion, we have set the magnitude of our spins, $\hbar S$, to 1. Therefore, $\hbar \sim 1/S$ and thus the energy of individual photons, $\hbar\omega(\vec{q})$, also scales as $1/S$. In our simulations, $S \to \infty$ and so $T \gg \hbar\omega$ at any nonzero temperature. Photon populations are thus always large and can be thought of as being in a coherent, classical state.

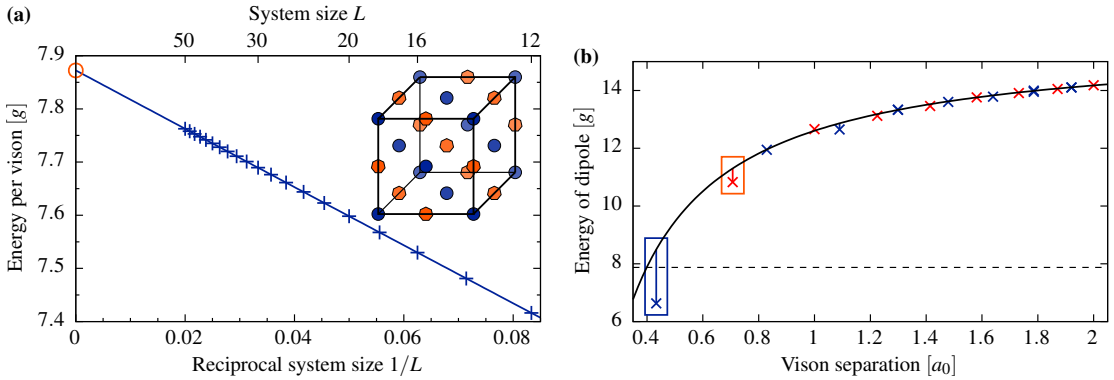

Figure 2.8. **(a)** Energy per vison in a rock-salt arrangement of visons **(inset)**, as a function of linear system size L. In a Coulomb interacting system, this energy is linear in $1/L$, see (2.22); the data indeed follow this trend excellently. The energy of an isolated vison is given by extrapolating to $L = \infty$ (red circle); the Coulomb interaction strength can be extracted from the slope of the fitted straight line. **(b)** Energy of a pair of nearby positive and negative visons as a function of separation. Energy estimates based on the full semiclassical theory match the Coulomb law $\varepsilon(r) = 2\mu - \alpha/r$ (solid line) well for all but nearest and second neighbours (boxed). In particular, the energy cost of a nearest-neighbour dipole is smaller than that of a single vison (dashed line).
Figures taken from Ref. 2.

b on the chosen plaquette by $-3\pi/2$ and by $\pi/2$ on its neighbours. Assuming that b was small initially, one can regard $\Delta b = -3\pi/2$ as changing b by $\pi/2$, together with a 2π phase slip that inserts a vison pair across the plaquette. That is, our "fundamental step" inserts a pair of visons with a symmetric near-field pattern around them; repeating it along a chain of sites amounts to moving the visons one apart from another, similar to inserting of spinons in classical spin ice.

The state generated by these local operations is, however, not the least energetic one, for it does not capture the $1/R^2$ far field of visons and contains a prominent "Dirac string" between them. While it may be possible to construct operators acting on all spins that have a larger overlap with the "true vison-creation operator" [15], we adopted a simpler and more straightforwardly reliable approach. We equilibrate the photons generated by the "fundamental step" described above using the Monte-Carlo algorithm and gradually reduce the effective temperature until the remaining photon population can be ignored, leaving behind a two-vison metastable state. The only issue with this method is vison movement: visons in the semiclassical model are not inherently mobile [97] (there are no explicit vison-hopping terms in the Hamiltonian), but a high-temperature photon cloud can move them around. We find that starting the annealing procedure at a sufficiently low temperature prevents such motion unless the visons are introduced within a distance of about twice the cubic lattice parameter.

The energy difference between the resulting configuration and the ground state can be regarded as the energy cost of two visons plus their Ewald-summed interaction energy. The latter, however, contains a surface term [127], which makes interpreting the results complicated if the vison arrangement has a net dipole moment. To remedy this problem, we used the rock-salt configuration shown in Fig. 2.8(a). This arrangement was set up using the vison insertion protocol described above; photons were equilibrated for $16L$ Monte-Carlo steps[†] at $T = g/256$, where the simulation box consists of $L \times L \times L$ cubic unit cells; then cooled 32 times by a factor of two and equilibrated for $8L$ steps each time. In the end, the temperature of the photon cloud reached $2^{-40}g \approx 9 \times 10^{-13}g$; equilibration was monitored through the acceptance rate of Metropolis steps, which remained consistently high throughout the process. The so obtained metastable energies are plotted in Fig. 2.8(a).

Modelling visons as Coulomb interacting charges, the energy of the rock salt configuration is

$$\varepsilon(L) = \mu - \frac{M\alpha}{La_0} \tag{2.22}$$

per vison, where μ is the energy cost of an isolated vison, a_0 is the cubic lattice parameter, α is the Coulomb interaction strength, and $M = 1.74756459\ldots$ is the appropriate Madelung constant. Fitting this form to the numerical results plotted in Fig. 2.8(a) yields

$$\mu = 7.872367608(68)\, g; \qquad \alpha = 3.1416145(37)\, ga_0. \tag{2.23}$$

The numerical results can be compared to an analytic estimate of the energy cost and interaction strength of visons from a quadratic approximation to the energy of the magnetic field, $-g\cos b$ (Appendix A.2). The quadratic estimate of μ, $8.858g$, is substantially larger than the numerical result (2.23); by contrast, the interaction strength, $\alpha = \pi g a_0$, matches excellently. This is so because the vison energy cost includes that of its immediate neighbourhood, where the quadratic approximation breaks down; on the contrary, most of the interaction energy is due to the overlap of far fields that are captured accurately.

We finally consider how accurately the interaction of nearby visons is described by the asymptotic Coulomb law. Since these visons were unstable against the photon cooling protocol, an alternative technique had to be used. Within quadratic approximation, the b-field configuration that minimises the magnetic energy for a given arrangement of visons can be found explicitly, as described in Appendix A.2: the sum of $-g\cos b$ for all

[†] A "Monte-Carlo step" in this work consists of the following: sampling the xy phase angle ϕ of each spin; a Metropolis attempt to change σ^z around each hexagonal plaquette; and sampling the gauge freedom of ϕ on each tetrahedron. These elementary steps are described in more detail in §2.3.3.

plaquettes provides a variational upper bound on the energy of the vison configuration. We benchmarked this estimate against the photon cooling technique at vison separations where the latter is applicable; because of the good agreement there, we decided to use this method to estimate the energy of vison dipoles with separations smaller than $2a_0$. The results are plotted in Fig. 2.8(b) together with the Coulomb energy estimate $2\mu - \alpha/r$, where μ and α are taken from (2.23). The latter is a remarkably good approximation even at third-neighbour distance; there is a discrepancy of about $0.5g$ for second neighbours and about $2g$ for nearest neighbours.

Importantly, the energy of a nearest-neighbour dipole is smaller than that of a single isolated vison, μ. Visons in QSI thus form a *weak electrolyte*. While deconfined, their dissociation is so energetically unfavourable that most visons at low temperatures remain associated with an oppositely charged one. This behaviour is quite different from that of spinons in CSI, whose energy cost is set independently by the dominant Ising exchange interaction; their dissociation is energetically favourable even at nearest-neighbour distance in the presence of entropic and dipolar Coulomb interactions.

2.4.3 Thermodynamics of visons and photons

Gapped excitations can naturally be introduced in numerical simulations via thermal fluctuations at finite temperature, which creates an equilibrium population of such excitations. Thermodynamic quantities like heat capacity or thermal conductivity are promising signatures of gapped quasiparticles and QSL behaviour in general [83, 99, 128]. Therefore, we studied the interactions of visons and photonic modes in a finite-temperature ensemble. This was greatly aided by the ability of our method to directly access observables such as the emergent magnetic field and the vison charge operator (2.21).

Magnetic-field pinch points

We performed static Monte-Carlo simulations on $20 \times 20 \times 20$ cubic unit cells of the pyrochlore lattice (128 000 spins) at 55 temperature points (distributed uniformly in $1/T$) between $0.4g$ and g. For temperatures above $0.5g$, 131 072 stochastically independent Monte-Carlo samples were generated; for those between $0.4g$ and $0.5g$, 262 144 samples were used. We evaluated static correlation functions of the emergent magnetic field, $\langle b(-\vec{q})b(\vec{q})\rangle$, and plotted them in the (hhk) plane for three temperatures in Fig. 2.9; correlators of the emergent electric field σ^z are also shown for comparison. At low temperatures, both correlators exhibit sharp pinch points at the Γ points in the pattern familiar from classical spin ice [12, 60], cf. (1.27). The electric-field pinch points remain sharp at all temperatures as spinons are excluded by the Monte-Carlo algorithm. On the contrary,

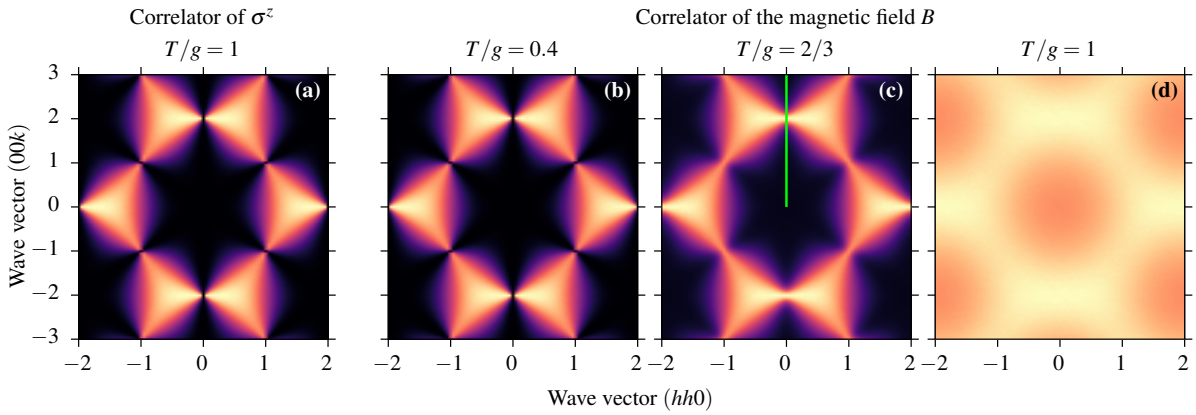

Figure 2.9. Static correlation functions $\langle \sigma^z(-\vec{q})\sigma^z(\vec{q})\rangle$ **(a)** and $\langle b(-\vec{q})b(\vec{q})\rangle$ **(b–d)** in the semiclassical QSI model. Due to the no-spinon constraint enforced in our Monte-Carlo algorithm, the former retains sharp pinch points at all temperatures. Unlike spin-1/2 QSI [19, 95], pinch points are not depleted near the Γ points. At low temperatures, the emergent magnetic field shows identical pinch points: these are blurred by thermally induced visons at higher temperatures, and at $T \gtrsim g$, they are washed out altogether. Figure taken from Ref. 2. Perceptionally uniform colour map chosen following Ref. 125.

the thermally introduced visons blur the b-field pinch points at finite temperatures in much the same way spinons blur CSI pinch points [66, 76]. This picture is in some departure from spin-1/2 QSI in which pinch points are suppressed near the Γ points (§2.1.4) and thus no sharp features appear [19, 95]. This is due to the different way in which photons appear in the two systems. In the semiclassical case, they are classical normal modes, hence their energy content is constant as per equipartition (cf. §2.4.1); in the spin-1/2 case, low-temperature physics is dominated by quantum zero-point fluctuations of photon modes which give rise to correlators proportional to the photon dispersion $\omega(\vec{q})$, thus suppressing the pinch points.

Pinch-point blurring is a common experimental diagnostic of spinons in CSI [66]; likewise, we were able to extract quantitative information about the visons from the blurring of b-field pinch points. We focus on the $(00k)$ axis (green line in Fig. 2.9), where the photon contribution to the correlator vanishes [12, 60], that is, the signal is entirely due to visons. Cuts of the correlator along this axis are plotted in Fig. 2.10(a) for four temperatures. These show an apparently Lorentzian peak at the (002) pinch point, indicating a Debye plasma of visons. Unlike the CSI case, however, our blurring pattern is not explained by a Lorentzian peak in itself: in particular, the peak appears on top of a substantial constant background [Fig. 2.10(b)]. This constant correlator can be ascribed to a large population of nearest-neighbour vison dipoles that are independent of the Debye plasma mentioned above (Appendix A.3). This underlines the observation that visons in

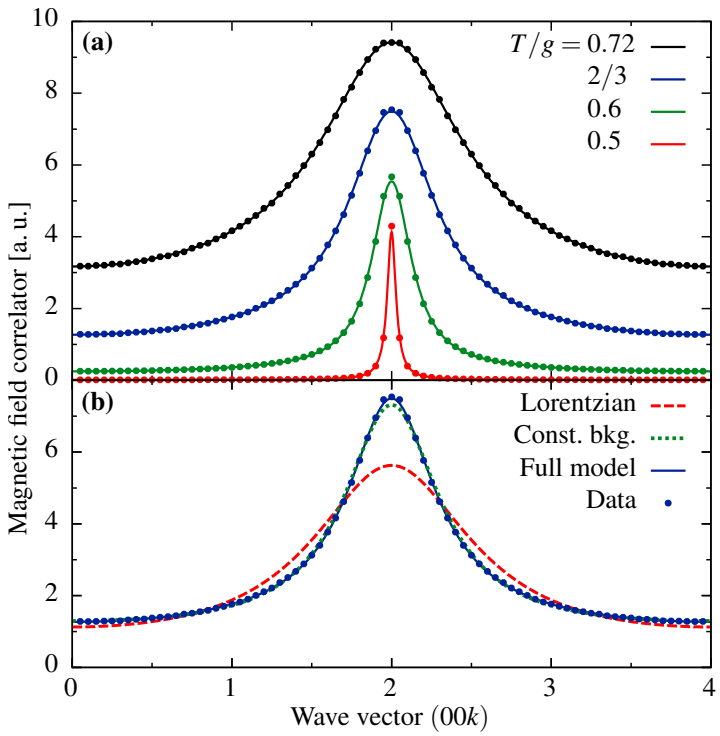

Figure 2.10. **(a)** Static correlation function of the emergent magnetic field, $\langle b(-\vec{q})b(\vec{q})\rangle$, along the (00$k$) axis for four different temperatures in the semiclassical QSI model. All data sets (coloured dots) fit excellently to the theoretical form (2.24) (solid lines).
(b) Simulation data at $T = 2g/3$ (blue dots) fitted to several theoretical models. A simple Lorentzian (red dashed line), corresponding to a Debye plasma of visons, provides a poor fit. An additional constant background, corresponding to nearest-neighbour vison dipoles, improves the fit substantially (green dotted line). A further minor, but significant, improvement can be achieved by including a contribution for second-neighbour vison pairs, giving rise to (2.24) (blue line).
Figures taken from Ref. 2.

QSI form a weak electrolyte, that is, they interact strongly enough that a large fraction of their thermal population remains associated, as discussed in §2.4.2.

Furthermore, the Coulomb energy formula for vison pairs substantially overestimates the energy cost of second-neighbour vison dipoles [Fig. 2.8(b)]. Therefore, we anticipate an excess population of them compared to that predicted by the Debye plasma approximation. While this effect is not qualitative, it does introduce a correction to the b-field correlator that is proportional to $\sin^2(q_z/8)$ along the (00k) axis (Appendix A.3.2). To take the effects of these closely associated dipoles into account, we fitted the func-

tional form

$$\langle b(-\vec{q})b(\vec{q})\rangle = C_1 + C_2 \sin^2(q_z/8) + \frac{C}{1 + \left[8\xi \cos(q_z/8)\right]^2} \tag{2.24}$$

to the data at all temperature points: $\xi \propto \rho_{\text{free}}^{-1/2}$ is the Debye length of the plasma formed by dissociated dipoles, where ρ_{free} is the density of dissociated visons; C_1 is proportional to the density of nearest-neighbour vison pairs, ρ_{nn}.[†] Equation (2.24) fits the data excellently throughout the temperature range we studied; Fig. 2.10(b) demonstrates that all three terms are necessary to achieve this, although very good agreement is reached even without the C_2 contribution.

Temperature dependence of vison density

Visons can also be counted explicitly in our Monte-Carlo simulations by evaluating (2.21) for a sampled spin configuration on all dual diamond sites. Besides the number N of all visons, the number N_{nn} of dual-diamond-lattice links with two oppositely charged visons on their ends was obtained; from these, the number of dissociated visons was estimated as $N_{\text{free}} = N - 2N_{\text{nn}}$. The density of bound dipoles and free visons is plotted as a function of inverse temperature in Fig. 2.11, together with the constant background C_1 of the b-field correlator and the Debye length ξ, respectively: we expect $\xi \propto \rho_{\text{free}}^{-1/2}$ [64] and $C_1 \propto \rho_{\text{nn}}$ (Appendix A.3.2). These relations hold quite well throughout the temperature range, which confirms that magnetic-field pinch-point blurring is a good measure of vison populations.

The densities of both dissociated and bound visons follow an approximate Arrhenius law at low temperatures with an apparent gap close to the bare vison energy (2.23) and the nearest-neighbour dipole energy shown in Fig. 2.8(b), respectively. They, however, saturate at $T \approx g$, an order of magnitude below the zero-temperature energy cost of visons; correspondingly, densities at $T \lesssim g$ are much larger than predicted by a simple Boltzmann factor with quasi-equilibrium gaps (red lines in Fig. 2.11).

Thermodynamic effects of quasiparticle interactions

To identify the origin of this excess density of visons, we considered two different estimates of their energy cost in the thermal ensemble: the dependence of energy on vison number, quantified by $\mathrm{d}\langle E\rangle/\mathrm{d}N$, and the slope of the Arrhenius plot $\log N$ versus $1/T$. Both quantities were evaluated using data recorded at single temperature points,

[†]The complicated form of the Lorentzian is to account for the periodicity of the data imposed by the lattice; C_2 is related to the *excess* density of second-neighbour dipoles.

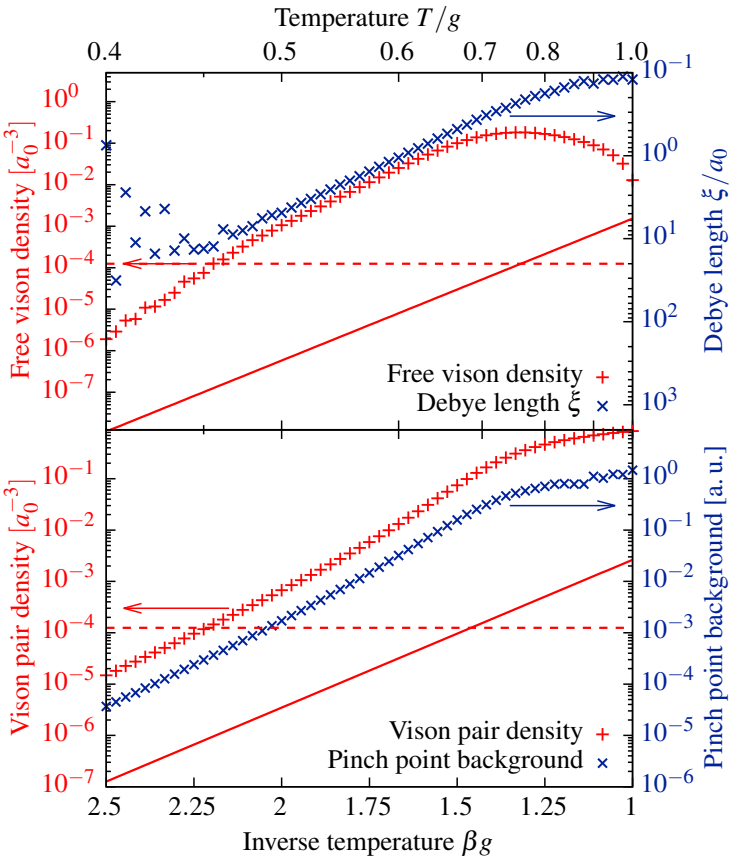

Figure 2.11. Arrhenius plots of the density of dissociated visons ρ_{free} (top panel) and nearest-neighbour vison dipoles ρ_{nn} (bottom panel), as well as the Debye length ξ (top panel) and background correlator C_1 (bottom panel) extracted from pinch point blurring data. The logarithmic scales of the left and right axes are linked to one another by the theoretical relations $\xi \propto \rho_{\text{free}}^{-1/2}$ and $C_1 \propto \rho_{\text{nn}}$, with an arbitrary scaling offset; the simulation data indeed follow these relations, even below the limit of a single vison (pair) in the entire system (red dashed lines). The densities of both vison populations follow an approximate Arrhenius relation with gaps similar to zero-temperature ones; their values, however, are much larger than a simple Boltzmann factor prediction (red solid lines). Figures taken from Ref. 2.

using the fluctuation–dissipation relations (Appendix A.4)

$$\mu_E = \left.\frac{\mathrm{d}\langle E\rangle}{\mathrm{d}N}\right|_T = \frac{\operatorname{cov}(E,N)}{\operatorname{var} N}, \qquad \mu_{\text{Arrh.}} = -\frac{\mathrm{d}\log N}{\mathrm{d}\beta} = \frac{\operatorname{cov}(E,N)}{N}. \tag{2.25}$$

These estimates are plotted in Fig. 2.12, together with $\mu_{\text{Arrh.}}/\mu_E = \operatorname{var} N/N$. At low temperatures, the distribution of N is generated by a Poisson distribution of thermal excitations which may well be a collective one made up of several visons: it follows (Appendix A.4) that $\operatorname{var} N/N$ gives the typical vison content of such a collective excitation.

At the lowest temperatures, $\operatorname{var} N/N$ tends to 2: this again demonstrates that the

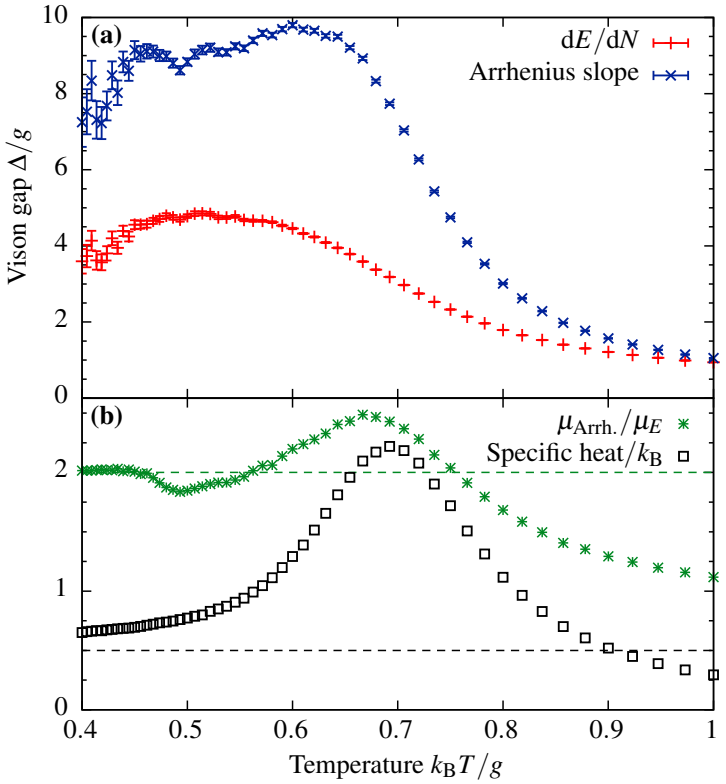

Figure 2.12. **(a)** Effective energy cost $\mu_E = \mathrm{d}\langle E\rangle/\mathrm{d}N$ (red) and effective Arrhenius gap $\mu_{\mathrm{Arrh.}} = -\mathrm{d}\log N/\mathrm{d}\beta$ (blue) of visons as a function of temperature. The latter is somewhat above the zero-temperature energy cost of nearest-neighbour vison dipoles at low temperatures. At $T \gtrsim 0.7g$, both gap estimates decline steeply.
(b) Ratio of the gap estimates, $\mu_{\mathrm{Arrh.}}/\mu_E$ (green stars), and specific heat (in units of k_B per spin) of the system (black squares) as a function of temperature. The former tends to 2 at low temperatures (green dashed line), as nearest-neighbour vison pairs are the dominant low-temperature vison population; at intermediate temperatures, larger coherent vison clusters raise it further. At low temperatures, the specific heat tends to $k_\mathrm{B}/2$ (black dashed line) due to equipartition of photon modes; it then increases due to photon interactions, turning into a vison Schottky peak at $T \approx 0.7g$; beyond that, the specific heat drops below the photon equipartition limit as photons and visons break down as useful quasiparticles.
Figures taken from Ref. 2.

dominant vison species at low temperatures are bound dipoles. At intermediate temperatures, $\mathrm{var}\, N/N$ increases further, to about 2.5, indicating collective excitations consisting of more than two visons. Salient examples of such clusters could be "polarons" consisting of a free vison and nearest-neighbour dipoles. As the b-field emanating from the former is quite large in its immediate neighbourhood, the energy cost of appropriately aligned dipoles decreases substantially, causing them to proliferate. In Appendix A.5, we construct a phenomenological model based on this mechanism and demonstrate that it

accounts for the qualitative features in Fig. 2.12. At even higher temperatures, var N/N decreases as N approaches saturation.

We also note that the energy cost of visons in the thermal ensemble declines steeply for $T \gtrsim 0.7g$. This is due to interactions between visons and photons arising from corrections to the gauge theory beyond quadratic order. The next order in the Villain expansion [17] consists of negative quartic terms, which, in the presence of a (thermal) population of photons, renormalise vison energies and interactions down. At intermediate temperatures, this renormalisation reduces the gap of visons compared to the zero-temperature value, resulting in larger thermal populations: this also explains the excess vison density observed in Fig. 2.11. We discuss this mechanism in more detail in Appendix A.5.

As temperature is increased, the energy cost of visons drops substantially and their density saturates. Eventually, at temperatures $T \gg g$, the visons cease to be useful quasiparticles as it is to be expected from the perspective of individual spins: the contribution of each hexagonal plaquette to the Hamiltonian (2.3) is $O(g)$, therefore, all spin configurations that satisfy the no-spinon constraint (which is still enforced by the Monte-Carlo algorithm) become roughly equally likely, regardless of vison content. The situation is similar to the crossover of classical spin ice into a high-temperature paramagnetic phase as spinons cease to be useful quasiparticles to describe the system. In our case, visons and photons are washed out at high temperatures, giving rise to a classical spin-ice phase with the spinon as its only excitation. Indeed, the heat capacity of the system, plotted in Fig. 2.12, displays a Schottky peak at $T \approx 0.7g$ above which it drops below the equipartition heat capacity of photon modes, $k_B/2$ per spin, and tends to zero as $T \to \infty$. This indicates that photons break down as quasiparticles together with the visons.

2.5 Outlook

To summarise, I developed a semiclassical numerical technique to simulate quantum spin liquids whose physics is underpinned by perturbative ring-exchange processes. Thus far, these systems remained elusive to large-S approaches, as their native bilinear exchange Hamiltonians order or form Ising spin liquids in this limit, devoid of QSL behaviour. Instead, our method uses effective ring-exchange Hamiltonians directly, which remain in the QSL phase in the large-S limit, allowing direct access to a variety of observables not readily available to other simulation methods. Our results provide detailed insight into the behaviour of quasiparticles in these quantum spin liquids, including gapped excitations that are not yet amenable to other computational and analytic tech-

niques.

I demonstrated the potential of our method on pyrochlore quantum spin ice, a paradigmatic U(1) QSL [15]. We identified a gapless, linearly dispersing branch of classical normal modes in spin dynamics. At low temperatures, these photonic modes are rather sharp and their dispersion exactly matches analytic results from large-S path integral calculations [17]. We also observed gapped, quantised vortices of the emergent magnetic field with the phenomenology expected for the elusive U(1) vison quasiparticles of QSI. In a showcase of the capabilities of our method, we were able to introduce these visons in a controlled way in the system and study their interactions in vacuum (that is, in an effectively zero-temperature photon background). Furthermore, we studied the interplay of photonic modes and visons in thermal equilibrium at finite temperature and found a clean signature of visons in the broadening of pinch points of relevant static correlators: a detailed analysis of these signatures uncovered substantial vison–vison and vison–photon interaction effects. While visons remain deconfined, these interactions cause them to associate in closely bound pairs, reminiscent of a weak electrolyte: this must be taken into account in modelling and understanding their behaviour, in particular in effective descriptions of vison dynamics [97, 129, 130].

One direct experimental signature that was identified in our work is a vison Schottky peak in the specific heat (Fig. 2.12). However, one must keep in mind that quantum photon excitations also contribute a Debye term to the specific heat, which is absent in semiclassical photonic normal modes. Hybridisation between photons and visons may well merge the two contributions and alter the shape of the Schottky anomaly. Additionally, the semiclassical numerical technique can naturally be applied to the dynamics of a high-temperature ensemble that contains thermally generated visons. However, the dynamics of semiclassical visons is far from trivial, since they are immobile at zero temperature [97] and their motion at finite temperature is due to being "tugged" by the photon background. Further work is needed to gain better insight into this behaviour.

Our method may also be extended to include terms in the simulated Hamiltonian that enable introducing static or dynamical spinons. Since spinons appear to have salient experimental signatures in, for instance, magnetisation and neutron-scattering measurements [98], a better understanding of spinon–vison interactions through our simulations may provide experimentally accessible handles to study collective photon and vison behaviour. It is important to note, however, that spinons are not quantised in the semiclassical setting (§2.3.3); therefore, understanding what their behaviour tells about the original quantum problem requires some care.

Finally, the ability of our technique to naturally include quantised charges [namely, U(1) visons] in an effectively classical system may have ramifications for lattice gauge

theories in general. Vison excitations in our semiclassical model are quantised solitons, similar to Dirac monopoles in quantum electrodynamics, and like those, they are likely to be in a strong-coupling regime. The reduction of vison energy by a thermally fluctuating background may indeed be a semiclassical analogue of running couplings that are brought about by virtual particle–antiparticle bubbles in QED. This is a speculative yet intriguing potential connection between QSI and the strong-coupling problem in QED, which warrants further investigation in future work.

3

Effective field theories for constrained quantum systems

Some of the exciting phenomena uncovered in strongly correlated systems in recent years – for instance, quantum topological order, deconfined quantum criticality, and emergent gauge symmetries – appear in systems in which the Hilbert space is effectively projected at low energies in a way that imposes local constraints on the original degrees of freedom. Cases in point include spin liquids, valence-bond systems, dimer models, and vertex models. Conventional field-theoretic approaches to such systems depend on physical insight to obtain a gauge-symmetric understanding of the constraint; they also lack generic quantitative predictive power for the coefficients of the terms that appear in their Lagrangians.

In this chapter, I devise a generalised route to obtain effective field theories for such systems using a slave-boson description coupled to a large-S path-integral formulation. I demonstrate the validity and capability of this approach by studying quantum dimer models (QDMs) and by comparing our results with the existing literature: the method overcomes systematic shortcomings of traditional field theories for QDMs and its results to leading order in the large-S expansion compare well with the known height description of the square-lattice QDM and the numerically estimated speed of light of the photon excitations on the diamond lattice. Finally, instanton considerations allow us to infer properties of the finite-temperature behaviour in two dimensions quantitatively.

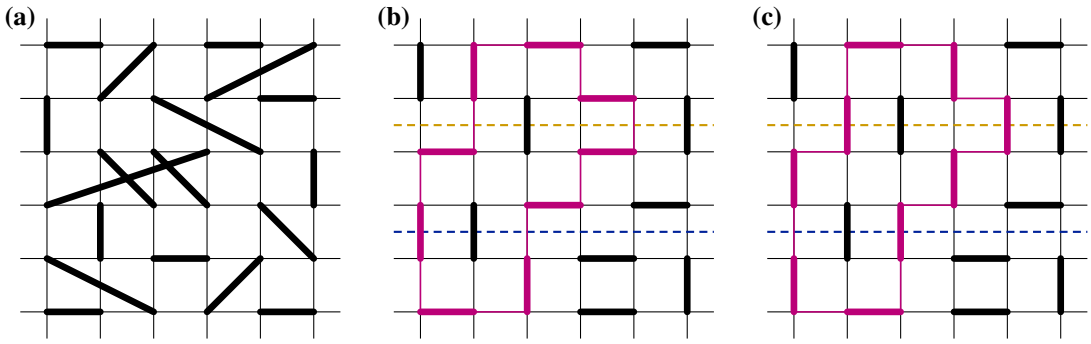

Figure 3.1. **(a)** Any $S = 0$ state of a spin-1/2 lattice model can be written as a superposition of product states of singlet dimers formed by arbitrary pairs of spins.
(b) In RVB states, dimers may only consist of nearest-neighbour sites.
(c) Due to the constraint that each lattice site belong to precisely one dimer, a valid dimer configuration can only be changed into another by flipping a closed loop (purple) of alternating dimers and empty links. On the square lattice, such a move preserves the number of dimers, taken with alternating signs, that intersect any horizontal or vertical line spanning the system: either a single dimer is moved to an empty link at an even distance (blue line), or two dimers at an odd distance are added/removed (gold line).

3.1 Classical dimer models

> *Now, since their natural form had been cut in two, each one longed for its own other half, and so they would throw their arms about each other, weaving themselves together, wanting to grow together.*
>
> — Plato, *Symposium*

Resonating-valence-bond (RVB) states have first been proposed by Anderson to describe the ground states of frustrated antiferromagnets on the triangle lattice [27, 28], as well as the fluctuating non-magnetic parent state of unconventional superconductors [29, 30]. Their construction begins with forming singlets ($S = 0$) out of nearby pairs of spin-1/2 electrons, thereby reducing their antiferromagnetic interaction energies under Heisenberg or Hubbard Hamiltonians. These singlets may form a regular pattern, resulting in ordered *valence-bond solid* states; an RVB state, however, is a superposition of all possible singlet coverings, which accounts for quantum fluctuations.

Allowing singlets formed out of any two spins [Fig. 3.1(a)], the set of singlet product states is indeed an (over)complete basis of the $S = 0$ sector of spin Hilbert space. In many important cases, however, nearest-neighbour singlets provide an accurate description of the physics already. Such short-range RVB states [Fig. 3.1(b)] have been studied extensively on their own right. Their connection to the problem of dimer coverings in

combinatorics and classical thermodynamics [9] allows us to gain insight into their exotic properties, including topological order and fractionalised excitations, purely from geometry, without detailed reference to their quantum dynamics [131]. This section is devoted to such *classical dimer models (CDMs).*

The constraint that each lattice site belong to precisely one dimer has a profound impact on the dynamics of dimer models, similar to the 2120 constraint on six-vertex models (§1.1.1). Dimers on the different links of the lattice are not independent degrees of freedom, and as such, they cannot be introduced into, or taken out of, a (classical) dimer configuration without violating the constraint. In particular, a valid dimer configuration can only be brought into another by identifying loops of alternating dimers and empty links, and flipping them [Fig. 3.1(b,c)], similar to loop updates in vertex models (§1.1). Given this correlated dynamics, it is not surprising that dimer coverings have *topological invariants,* which cannot be changed through local moves of this type.

I shall first demonstrate how fractionalised excitations can be constructed on top of the space of dimer configurations by breaking up single dimers. In the rest of the section, I discuss the topological invariants arising on bipartite and non-bipartite lattices, and their effects on these excitations.

3.1.1 Fractionalised excitations

The elementary spinful excitations of ordered (e.g., Néel) antiferromagnets are single-site spin flips: these may delocalise as spin waves, but they can always be described as superpositions of spin flips on single sites. By contrast, the simplest spin-1 excitation of an RVB basis state is created by turning a singlet dimer into a triplet. In the dimer language, this amounts to removing a dimer and replacing it with two spin-1/2 *monomers* [Fig. 3.2(a)]. Now, these monomers can be separated from one another by rearranging the dimer covering between them [Fig. 3.2(b)]; in the spin language, the triplet dimer is made long-ranged (but not the singlet ones).

If the dimer covering has some long-range order, the rearrangement necessary to separate the monomers violates the same, leading to an energy cost proportional to their distance [Fig. 3.2(c)]. In this case, the monomers are said to be *confined:* they are not true elementary excitations of the system, since they always appear in closely bound pairs, which can be coarse-grained into an ordinary spin flip.

By contrast, in an RVB state that includes all dimer coverings on an equal footing, the monomers can be separated without significantly changing the arrangement of the remaining dimers [Fig. 3.2(b)]. Indeed, the monomers could be recombined somewhere else in the lattice, leading to another valid dimer covering in the same RVB state. There-

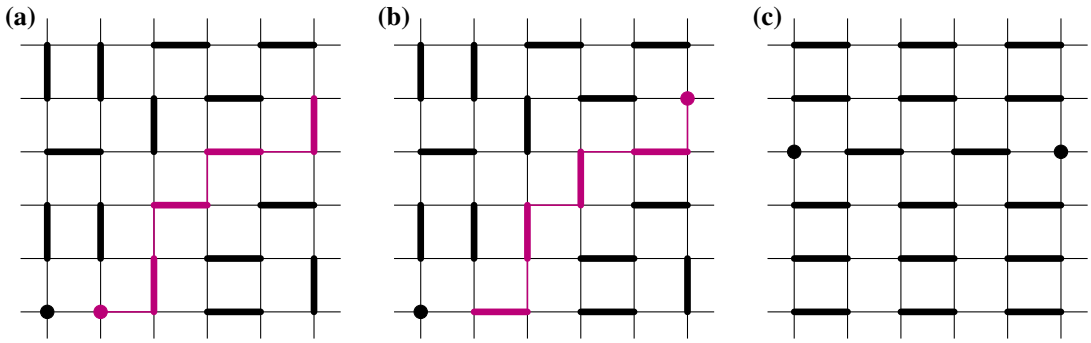

Figure 3.2. **(a)** Spin-1 excitations on top of an RVB state can be constructed by replacing a singlet dimer with a triplet one (shown as two monomers).
(b) The monomers can be separated by shifting other dimers around (purple). Away from the monomers, the dimer configuration remains consistent with the RVB state: the monomers become deconfined elementary excitations.
(c) In an ordered valence-bond-solid state, separating monomers breaks the ordering, incurring an energy cost proportional to the separation. These monomers remain confined.

fore, monomers in such a state are *deconfined* and behave as elementary spin-1/2 excitations of the system. As physical excitations always carry integer angular momentum, the monomers (also called spinons, cf. §1.1.1 and §1.2) are also *fractionalised*.

3.1.2 Bipartite lattices: winding numbers, height mapping, Coulomb phase

As a first example, consider the CDM on the 2D square lattice and focus on the dimers that intersect a horizontal (or vertical) line spanning the entire lattice (dashed lines in Fig. 3.1). Any closed loop on the lattice will cross such a line an even number of times: without loss of generality, let us consider a flippable loop (i.e., one containing alternating dimers and empty links) with two crossings. If the distance between these crossings is even (odd), there must be an even (odd) number of links in between them, regardless of the shape of the loop:

- in the first case, precisely one of the links that cross the line will have a dimer on it, both before and after the flip (blue dashed line in Fig. 3.1);
- in the second case, either both links were occupied and become empty upon flipping, or vice versa (gold dashed line).

In either case, the sum of dimer occupation numbers on all links that cross the line, *added with alternating signs,* remains constant, making it a topological invariant.

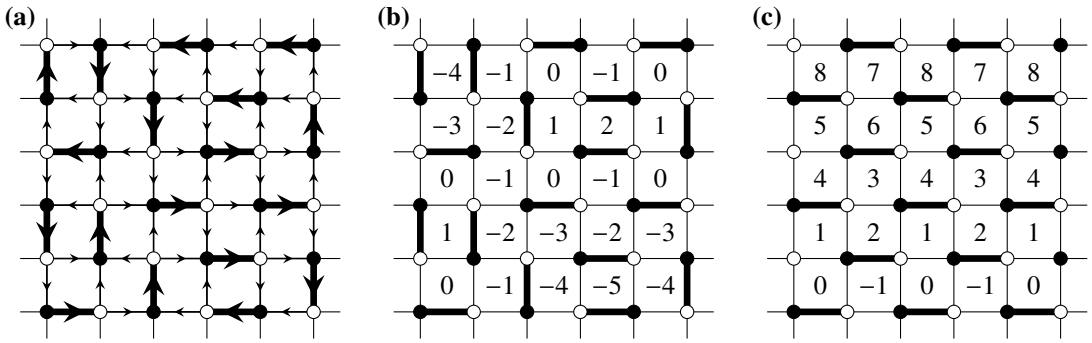

Figure 3.3. **(a)** Flux mapping on the square lattice. Each dimer is assigned flux $z - 1 = 3$ flowing from sublattice A (black dots) to B (white dots); empty links carry flux 1 from B to A. As a result, each lattice site has zero net flux.
(b) Height mapping on the square lattice. Each plaquette is given an integer height that changes by $z - 1 = 3$ or 1 upon crossing a dimer or an empty link. The lattice curl of this height field reproduces the flux field.
(c) In a staggered ordered state, the height field increases by four every two rows or columns, leading to the largest possible coarse-grained ∇h.

In open boundary conditions, this invariant is trivial: no updates can ever change it, so it remains zero for all lines. In periodic boundary conditions, however, one can also flip loops that wind around the entire lattice: these cross each such line an odd number of times, thus allowing the invariant to be changed by ± 1. It is important to point out that

(i) the invariant is still constant for *local* updates that do not wind around the periodic boundary conditions;

(ii) winding updates change it by the same amount on each horizontal or vertical line (up to signs that can be removed by defining the invariant properly). It follows that there are only two independent topological invariants, one each for lines that wind the lattice horizontally or vertically.

Due to the relationship between these invariants and winding loops, the former are usually called *winding numbers* and labelled W_x, W_y. They are integers (arbitrarily large in the thermodynamic limit, otherwise limited by system size) and define an extensive number of *topological sectors* that are only connected to one another by loop updates that wind around the periodic boundary conditions. The fact that these sectors are labelled by pairs of integers is a hallmark of *U(1) topological order*.

Two-dimensional bipartite lattices (i.e., ones that can be divided into two sublattices, A and B, without any links within a sublattice) also admit two convenient mappings of the dimer coverings, which allow us to study them using field-theoretical concepts [132]:

- Assign a flux of 1, flowing from sublattice B to A, to all empty links, and a flux of $z-1$, flowing from A to B, to all dimers, where z is the coordination number of the lattice [Fig. 3.3(a)]. Since there are $z-1$ empty links around each node, there is no net flux flowing into or out of them, making these fluxes a divergence-free lattice vector field (cf. §1.1.1).

- Assign a height h to all plaquettes of the lattice such that going around a node in sublattice A anticlockwise (or clockwise around one in sublattice B), h increases by $z-1$ when crossing a dimer, and decreases by 1 otherwise [Fig. 3.3(b)]. This can be done self-consistently since there are $z-1$ empty links and one dimer around each node, leading to no change in h upon fully encompassing a node.

The two mappings are related to one another: h can be thought of as a 2D vector potential (equivalent to the z component of a true 3D vector potential), whose lattice curl gives the divergence-free flux field.

Using the construction (1.4) introduced for the six-vertex model, the flux field can be coarse-grained into a *continuum* vector field $\vec{B}$; the vanishing lattice divergence of the fluxes implies that $\nabla \cdot \vec{B} = 0$. The arguments leading to (1.5) also apply, giving rise to the coarse-grained entropy functional

$$S = -\frac{\kappa}{2} \int \mathrm{d}^2 r |\vec{B}|^2. \tag{3.1}$$

That is, dimer models on bipartite lattices are Coulomb phases [14].

Monomers in the flux language become sources and sinks of flux: namely, the lattice divergence of the flux field around a monomer on sublattice A (B) is $-z$ ($+z$); after coarse-graining, this implies $\nabla \cdot \vec{B} = \pm z$. That is, analogous to spinons in the six-vertex model, monomers become quantised charges in the Coulomb phase (3.1), with entropic Coulomb interactions that scale as $\log R$ in 2D.

The winding numbers W_x and W_y give the net vertical and horizontal fluxes across loops that wind around the lattice: the divergence-free constraint implies that they only depend on the winding topology of the loop around the lattice, so they can be defined straightforwardly for non-square lattices. After coarse-graining, the same winding numbers give rise to a uniform background $\vec{B}$ field on the order of $(W_x, W_y)/L$, where L is the linear size of the lattice: in the thermodynamic limit, changing W by one is insignificant, but a finite winding number *density* can lead to observable physical differences, such as a reduction in entropy.

The discussion above extends straightforwardly to 3D lattices: the flux mapping and its coarse-graining remains identical, leading to a 3D Coulomb phase (3.1), with $1/R$

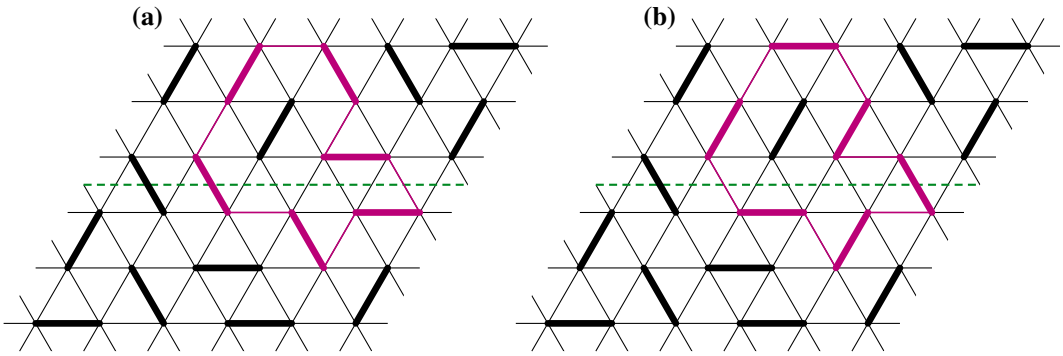

Figure 3.4. On a non-bipartite lattice, flipping a loop (purple) can change the number of dimers that intersect a system-spanning line (dashed green line), or any linear combination of the occupation numbers along it. The *parity* of intersecting dimers, however, is unchanged by local updates, providing two $\mathbb{Z}_2$ topological invariants.

Coulomb interactions between monomers [12, 14]. The height mapping is replaced by a full 3D vector potential whose lattice curl recovers the flux mapping.

3.1.3 Non-bipartite lattices: $\mathbb{Z}_2$ topological order

The flux and height mappings introduced above require the lattice to be bipartite in order to orient links and choose the sense of rotation around sites, respectively. Accordingly, on a non-bipartite lattice, no convention for link orientations results in a flux mapping that is either divergence-free around any given site, or would coarse-grain into a smooth $\vec{B}$ field. Without the divergence-free condition, no long-ranged Coulombic correlations arise: indeed, numerical studies show that dimer–dimer correlations decay exponentially with a correlation length ξ on the order of the lattice spacing for the CDM on both the triangular [133] and FCC [14] lattices. Accordingly, any entropic interaction between monomers is short-range, too, decaying over the same length scale.

Since the winding numbers $W_{x,y}$ on bipartite lattices can be thought of as the flux field linking the periodic boundary conditions, they cannot be consistently defined on non-bipartite lattices, either. Nevertheless, dimer coverings on these lattices still have topological invariants: since local loops (i.e., ones that do not wind around the periodic boundary conditions) still intersect any system-spanning line an even number of times, swapping dimers and empty links along them leaves the *parity* of dimers that intersect the same line unchanged[†] (Fig. 3.4). It is, however, changed by loops that wind around the system, consistently for all parallel lines: therefore, dimer coverings can be labelled

[†]This is equally true on bipartite lattices: there, the winding parities are the parities of the corresponding winding numbers.

using two *winding parities,* one each for lines that wind horizontally and vertically, analogous to the toric code (§1.1.3).

Such topologically protected parities are the hallmark of a $\mathbb{Z}_2$ *topological order* with four (or 2^d in d dimensions) distinct topological sectors. Unlike U(1) systems, where a large winding number can be detected as a uniform background flux, the short-ranged correlations imply that differences in local observables between RVB states in these sectors are $O(e^{-L/\xi})$, vanishing exponentially, rather than algebraically, as the system size increases. Accordingly, the corresponding QDMs have four nearly degenerate ground states that cannot be distinguished by local measurements, and are thus robust to local noise: such systems are useful for storing quantum information as *topologically protected qubits* [21, 26].

3.2 Rokhsar–Kivelson (RK) model

> *The formulation of a problem is often more essential than its solution.*
>
> — Albert Einstein, *The Evolution of Physics*

In order to endue an RVB model with quantum dynamics, it is natural to project an underlying antiferromagnetic Heisenberg Hamiltonian onto the subspace of nearest-neighbour dimer coverings [26, 89].[†] Similar to dimer moves in classical dimer models (Fig. 3.1), any offdiagonal matrix element of the resulting *quantum dimer model (QDM)* must connect states in the dimer basis via loop updates. In the spin language, these updates involve at least as many spin operators as the length of the loop, that is, they only arise perturbatively from the two-spin Heisenberg interactions. The perturbative expansion is dominated by the shortest loops[‡] along which a valid flip is possible: these consist of four links on the square, triangular, and cubic lattices, and six links on the honeycomb and diamond lattices; the corresponding *plaquette-flipping terms* are fourth and sixth order in the perturbation expansion, respectively.

The same perturbative expansion also generates diagonal terms; to the lowest orders, however, these are equal for all dimer coverings. The first nontrivial contributions are

[†]This projection is different from the usual Schrieffer–Wolff prescription as all terms of the Hamiltonian are of the same magnitude and dimer coverings (understood as spin states) are not orthogonal. Indeed, the "small parameter" in the perturbative expansion is the overlap of dimer coverings. On the other hand, the dimer states in the final RK Hamiltonian are treated as orthogonal [131]: the difference is corrected for in the process of the expansion expansion [26, 89].

[‡]These loops are usually the plaquettes of the lattice, except on the triangular lattice, where the elementary plaquette has three links, so it cannot be flipped. There, the RK model is defined in terms of rhombuses made up of two neighbouring triangles [26, 134]. Nevertheless, these operations are still referred to as plaquette flipping.

due to processes that flip a plaquette back and forth, resulting in an energy offset proportional to the number of *flippable plaquettes* in the covering. These processes are eighth order in the perturbative expansion on the square [89] and triangular [26] lattices; however, as the "small parameters" in the expansion are in fact order one, both diagonal and offdiagonal terms are important to give a good account of RVB states, and their relative strengths does not follow straightforwardly from the perturbation expansion. This motivates the *Rokhsar–Kivelson (RK) model* [89], where the coefficients of plaquette-flipping and counting terms, J and V, are free parameters:

$$H = \sum_p \Big[-J\big(|\square\rangle\langle\square| + |\square\rangle\langle\square|\big) + V\big(|\square\rangle\langle\square| + |\square\rangle\langle\square|\big)\Big]. \tag{3.2}$$

$|\square\rangle$ and $|\square\rangle$[†] are shorthand for all pairs of dimer coverings that only differ by having two horizontal or two vertical dimers on plaquette p, respectively. Indeed, the first term of (3.2) flips the plaquette p, while the second results in a diagonal energy term if p is flippable. Without loss of generality, J can be taken as positive [26, 89].

An important feature of the RK model (3.2) is the existence of the exactly solvable *RK point* at $V = J$, where the Hamiltonian can be rewritten as

$$H = J \sum_p \big(|\square\rangle - |\square\rangle\big)\big(\langle\square| - \langle\square|\big), \tag{3.3}$$

a sum of dyads with positive prefactor J. As (3.3) is manifestly positive semidefinite, a state with zero energy would be its ground state. Now, an equal-amplitude superposition of *all dimer coverings,*

$$|\text{RVB}\rangle = \sum_c^{\substack{\text{topo.}\\ \text{sector}}} |c\rangle, \tag{3.4}$$

has an equal overlap with $|\square\rangle$ and $|\square\rangle$ in every term of (3.3), therefore,

$$\big(\langle\square| - \langle\square|\big)|\text{RVB}\rangle = 0 \implies H|\text{RVB}\rangle = 0.$$

That is, the RVB liquid state (3.4) is the ground state of (3.3). The expectation value of any observable that is diagonal in the dimer basis is the same for $|\text{RVB}\rangle$ as for the classical dimer liquid discussed in §3.1; properties of excited states can also be inferred from classical Langevin dynamics of the latter [135]. On bipartite lattices, the ground state of (3.3) shows algebraic dimer correlations and a quadratically dispersing gapless excitation spectrum [89, 135, 136]; on non-bipartite lattices, dimer correlations are short-

[†]This graphical representation is consistent with the square lattice; equivalent terms can equally well be written down for any other lattice, see for instance (2.7).

range and excitations are gapped [136, 137]. The $|\text{RVB}\rangle$ states also inherit the U(1) or $\mathbb{Z}_2$ topological structure of the classical liquids; indeed, as all plaquette-flipping terms in (3.2) are local, topological sectors remain disconnected, resulting in a zero-energy ground state in each sector.

For $V > J$, a similar restructuring of the Hamiltonian shows that it is positive definite for states that have any overlap with any of the $|\square\rangle$ or $|\square\rangle$; by contrast, *staggered* states, which have no flippable plaquettes at all, will still have zero energy, and thus become the ground state. On the other hand, as $V \to -\infty$, flippable plaquettes become energetically favourable, leading to *columnar ordered* states that maximise their number (Fig. 3.5). For intermediate values of V, several other ordered states may form with various patterns of resonating plaquettes. While these depend on details of the lattice, the fate of the RVB liquid phase away from the RK point can be predicted from field-theoretic arguments that only depend on dimensionality and bipartiteness, as discussed in the next section.

3.2.1 Effective gauge theories

Classical dimer coverings on a 2D bipartite lattice are captured by the height mapping, which can be coarse-grained and used as the starting point of a quantum-field-theoretic description. Analogous to Landau theory, an effective action, consisting of low-order derivatives of h, can be written down [131, 138, 139]:

$$S[h] = \int \mathrm{d}\tau\, \mathrm{d}^2 r \left[\frac{1}{2}(\partial_\tau h)^2 + \frac{c^2}{2}(\nabla h)^2 + \frac{K^2}{2}(\nabla^2 h)^2 - \lambda \cos(2\pi h/z) \right]; \tag{3.5}$$

h^2 terms are absent as the physics is invariant under shifting h by a constant. Before coarse-graining, h is an integer, and its modulus with respect to the coordination number z on each plaquette is fixed for all dimer coverings: the last term in (3.5) enforces these constraints after coarse-graining [131, 138]. The action (3.5) almost always describes ordered states:

- If $c^2 < 0$, the action is minimised if ∇h is as large as possible. In the fine-grained image, ∇h is maximised by a staggered order [Fig. 3.3(c)]; therefore, this regime is consistent with $V > J$.

- If $c^2 > 0$, the action is the same as that of compact quantum electrodynamics (QED) in 2D. In this case, the cosine term is a relevant perturbation to the non-compact theory given by the first two terms of (3.5), leading to a proliferation of instantons that gap out the photons, resulting in an ordered state [16] (see §3.4.3 for a detailed discussion of this mechanism).

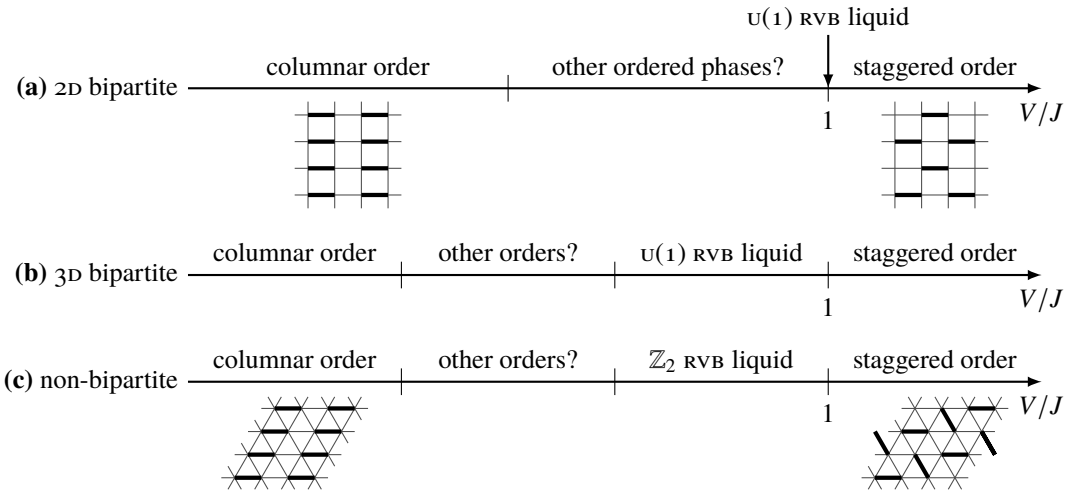

Figure 3.5. Generic phase diagram of the RK model on different lattices as a function of the dimensionless parameter V/J [131]. Above the RK point, all models have a staggered ground state with no flippable plaquettes; for sufficiently negative V/J, columnar orders form with as many flippable plaquettes as possible.
(a) On 2D bipartite lattices, the U(1) RVB state (3.4) is restricted to the RK point. At $V/J \lesssim 1$, instanton effects gap out the emergent photon, resulting in an ordered phase [16].
(b) On 3D bipartite lattices, instanton effects are irrelevant, so the U(1) liquid phase remains stable for an extended range of $V/J < 1$, before giving way to ordered states. The ordered states are illustrated on a cubic lattice in Fig. 3.10.
(c) The $\mathbb{Z}_2$ RVB liquid state that forms at the RK point on non-bipartite lattices is gapped, so it is protected over a finite range of $V/J < 1$.
In all cases, the presence of ordered phases other than the columnar one depends on lattice details; various resonating plaquette phases are illustrated in Figs. 3.8, 3.10, and 3.11.

- At $c^2 = 0$, the free field theory consists of the $(\nabla^2 h)^2$ and $(\partial_\tau h)^2$ terms, for which the cosine perturbation becomes irrelevant. The excitations of this theory are gapless with a quadratic dispersion $\omega = Kk^2$ [89, 135, 136]; correlations in h decay algebraically, consistent with the RK point.

These observations support (3.5) with $c^2 \propto (J - V)$ as an effective field theory for the RK model on bipartite 2D lattices. It follows that the RVB liquid (3.4) is not part of a wider phase but appears only at the RK point, giving rise to the generic phase diagram in Fig. 3.5(a).

Effective theories for 3D bipartite lattices can be constructed using similar arguments in terms of the coarse-grained flux field $\vec{B}$ and a vector potential $\vec{A}$ such that $\vec{B} = \nabla \times \vec{A}$. Keeping the dominant terms, we obtain [12, 15, 140]

$$S[\vec{A}] = \int \mathrm{d}\tau \, \mathrm{d}^3 r \left[\frac{1}{2} (\partial_\tau \vec{A})^2 + \frac{c^2}{2} \vec{B}^2 + \frac{K^2}{2} (\nabla \times \vec{B})^2 \right], \tag{3.6}$$

where $c^2 \propto (J - V)$ again. The first two terms describe Maxwellian electromagnetism with linearly dispersing photon excitations. Unlike in 2D, the compact nature of the theory [which would appear as additional cosine terms in the action (3.6)] is irrelevant [16], therefore, this U(1) RVB phase survives for finite range of $V/J \leq 1$ [Fig. 3.5(b)]. At the RK point, the second term in (3.6) vanishes, leading to quadratically dispersing photons [136].

Finally, on non-bipartite lattices, the $\mathbb{Z}_2$ topological order observed in §3.1.3 suggests that the excitations of the corresponding RK models can be described in terms of the $\mathbb{Z}_2$ lattice gauge theory discussed in §1.1.3; monomers play the role of gauge charges, i.e., spinons. Indeed, the spectrum of the triangular lattice model at the RK point is gapped [136, 137]: this gap protects the topologically ordered RVB liquid state for a finite range of $V/J \leq 1$ [Fig. 3.5(c)]. By construction, monomers in the RK model are non-dynamical; the lowest-lying excitations [137] are visons, which live on the triangular plaquettes of the lattice; a variational wave function that contains a pair of them on plaquettes A and B is given by [137, 141, 142]

$$|AB\rangle = (-1)^{\hat{N}_{AB}} |\text{RVB}\rangle. \tag{3.7}$$

Here, $\hat{N}_{AB}$ is the operator that counts the number of dimers intersecting an arbitrary path between the plaquettes A and B: the choice of path is immaterial as the *parity*[†] of this number is independent of it. Carrying a monomer around any loop that encloses a single vison changes the occupation of precisely one link along this path, introducing a factor of -1: that is, monomers and visons obey mutual semionic statistics, as expected of the gapped excitations of a $\mathbb{Z}_2$ gauge theory (§1.1.3). Visons remain the lowest-lying excitations of the triangular lattice RK model away from the RK point as well; their condensation at $V/J \approx 0.8$ causes an ordering transition out of the liquid phase [143].

[†]On a bipartite lattice, the total flux of $\vec{B}$ crossing such a path would be invariant, allowing any unit complex number to be used in (3.7) instead of -1: this would result in continuous insertion of the U(1) flux dual to $\vec{B}$.

3.3 Field theories via soft constraints

A soft answer turneth away wrath: but grievous words stir up anger.

— Proverbs 15:1

We represent quantum dimer models using a secondary Hilbert space in which we assign a bosonic mode $\hat{b}_\ell$ to each link ℓ of the lattice [144]. We associate the number of dimers on a link with the occupation number of the corresponding mode, thus embedding the dimer Hilbert space in the larger Hilbert space of the bosons. The constraint that each site of the lattice belong to precisely one dimer can be expressed in the boson language as

$$\Pi_r \equiv \sum_{\ell \in \vec{r}} \hat{b}_\ell^\dagger \hat{b}_\ell - 1 = 0, \tag{3.8}$$

where the sum runs over all links attached to the vertex at position $\vec{r}$. This constraint implies that the bosons are hard-core: $\hat{n}_\ell \equiv \hat{b}_\ell^\dagger \hat{b}_\ell = 0, 1$.

Any QDM Hamiltonian admits several such representations, as they only need to agree for the small segment of the boson Hilbert space that obeys (3.8). In particular, the simplest representation of the RK Hamiltonian (3.2) with plaquettes of four links[†] is

$$H_D = \sum_p \left[-J \left(\hat{b}_1^\dagger \hat{b}_3^\dagger \hat{b}_2 \hat{b}_4 + \text{H.c.} \right) + V \left(\hat{b}_1^\dagger \hat{b}_1 \hat{b}_3^\dagger \hat{b}_3 + \hat{b}_2^\dagger \hat{b}_2 \hat{b}_4^\dagger \hat{b}_4 \right) \right], \tag{3.9}$$

where the sum runs over the plaquettes p and the indices $1, \ldots, 4$ refer to links around the same plaquette in order. We note that (3.9) is homogeneous in the operators $\hat{b}$: as a result, it can be recast similarly to (3.3) at the RK point, allowing an exact RVB ground state to be constructed for arbitrary S (Appendix B.1).

3.3.1 Path integral formulation and large-S approximation

Using the (overcomplete) basis of bosonic coherent states,

$$|\beta\rangle = \exp(\beta \hat{b}^\dagger - \overline{\beta} \hat{b}) |\text{VAC}\rangle = e^{-|\beta|^2/2} \sum_{n=0}^{\infty} \frac{\beta^n}{\sqrt{n!}} |n\rangle_b, \tag{3.10}$$

[†]Hamiltonians for plaquettes of different sizes can be constructed analogously, see §3.5.2.

the partition function of the Hamiltonian (3.9) can be represented in terms of the coherent state path integral

$$\mathcal{Z} = \int \mathcal{D}(\overline{\beta}, \beta) e^{-\mathcal{S}[\overline{\beta},\beta]}; \qquad \mathcal{S} = \int \mathrm{d}\tau \left[\sum_{\ell} \overline{\beta}_\ell \partial_\tau \beta_\ell + H_D[\overline{\beta}, \beta] \right], \tag{3.11}$$

where the ladder operators $\hat{b}_\ell^\dagger$, $\hat{b}_\ell$ are replaced by complex-valued functions $\overline{\beta}_\ell(\tau)$, $\beta_\ell(\tau)$, over which the path integral is performed [145]. In this formalism, the hard dimer constraint (3.8) can be enforced exactly via additional path integrals over Lagrange-multiplier fields $\lambda_r(\tau)$ for each lattice site r that are coupled to the conserved Π_r. Performing these integrals, however, gives rise to a range of additional interactions between the variables β, rendering the problem intractable.

Instead, we replace (3.8) with a *soft constraint* by fixing the sum of $|\beta_\ell|^2$, rather than that of $\hat{b}_\ell^\dagger \hat{b}_\ell$, on the links that surround each vertex. This is an approximation because the coherent states (3.10), represented by the variables β_ℓ, have no well-defined boson number; its quality, however, improves with increased boson number, as $\mathrm{var}\, \hat{n}_\ell = \langle \hat{n}_\ell \rangle = |\beta_\ell|^2$ in a coherent state, so relative number fluctuations decay as $1/\sqrt{n_\ell}$. Therefore, the soft constraint is physically equivalent to a large-S representation, in which the hard constraint (3.8) is replaced by

$$\Pi_r \equiv \sum_{\ell \in v_r} \hat{b}_\ell^\dagger \hat{b}_\ell - S = 0 \quad \Longleftrightarrow \quad \sum_{\ell \in v_r} |\beta_\ell|^2 = S \tag{3.12}$$

for some $S \gg 1$, without changing the form of the Hamiltonian (3.9).

3.3.2 Gaussian approximation

In what follows, it is convenient to introduce the radial gauge

$$\beta_\ell = \sqrt{\rho_\ell}\, e^{i\Phi_\ell} \equiv \sqrt{\frac{S}{z} + \delta\rho_\ell}\, e^{i\Phi_\ell} \tag{3.13}$$

for the bosonic fields, where z is the coordination number of the lattice. The latter form is useful in a dimer-liquid phase, where the typical occupation number ρ of each link is equal up to small fluctuations $\delta\rho$. In this representation, the Berry phase $\sum_\ell \overline{\beta}_\ell \partial_\tau \beta_\ell$ gives rise to the term $\sum_\ell i\, \delta\rho_\ell\, \partial_\tau \Phi_\ell$, as well as total-derivative terms that do not contribute to the action. On bipartite lattices, the canonically conjugate "number operator" $\delta\rho$ and angular variable Φ are analogous to the electric field and magnetic vector potential of a compact U(1) gauge theory, respectively (§1.1.2, §2.1.1): while coarse-graining the flux field into a divergence-free *magnetic* field, as done in §3.1, is convenient for building ad

hoc field theories, it is more natural to consider monomers to be gauge charges that are divergences of the *electric* field (§2.1.3).

The Hamiltonian (3.9) and the soft constraint (3.12) can be rewritten as

$$H_D = \sum_p \left[-2J\sqrt{\rho_1\rho_2\rho_3\rho_4}\cos(\Phi_1+\Phi_3-\Phi_2-\Phi_4)+V\rho_1\rho_3+V\rho_2\rho_4\right] \tag{3.14}$$

$$0 = \sum_{\ell\in v_r}\delta\rho_\ell. \tag{3.15}$$

We shall define $\phi_p \equiv \Phi_1+\Phi_3-\Phi_2-\Phi_4$. The Hamiltonian (3.14) is highly nonlinear in the parameters ρ and Φ; as a starting point, we shall expand it to quadratic order around its minima. In the dimer-liquid phase, the energy is lowest at $\rho_\ell \equiv S/z$, therefore, the Gaussian approximation follows from taking $\delta\rho \ll S/z$. Furthermore, for a given set of ρ, (3.14) is minimised if $\phi_p = 2\pi n$ ($n \in \mathbb{Z}$) for each plaquette p.

To expand in $\delta\rho$, let us first rewrite $\cos\phi = 1-(1-\cos\phi)$ and notice that the term in square brackets contains only quadratic and higher-order contributions. Therefore, the square root in the second term of

$$\sqrt{\rho\rho\rho\rho}\cos\phi = \sqrt{\rho\rho\rho\rho} - \sqrt{\rho\rho\rho\rho}\,(1-\cos\phi)\,, \tag{3.16}$$

needs to be expanded only to leading order in S: $\sqrt{\rho\rho\rho\rho} \simeq (S/z)^2$. Upon expanding the first term, as well as the potential terms of (3.14), one obtains both linear and quadratic terms in $\delta\rho$. The former vanish upon summing over all plaquettes due to the dimer constraint (3.15). The constant and quadratic terms are given by

$$\sqrt{\rho\rho\rho\rho} \simeq \frac{S^2}{z^2} + \frac{1}{4}(\delta\rho_1\delta\rho_2+\delta\rho_1\delta\rho_3+\delta\rho_1\delta\rho_4+\delta\rho_2\delta\rho_3+\delta\rho_2\delta\rho_4+\delta\rho_3\delta\rho_4); \tag{3.17a}$$

$$-\frac{1}{8}(\delta\rho_1^2+\delta\rho_2^2+\delta\rho_3^2+\delta\rho_4^2)$$

$$\rho\rho+\rho\rho \simeq \frac{2S^2}{z^2} + \delta\rho_1\delta\rho_3+\delta\rho_2\delta\rho_4; \tag{3.17b}$$

the constant terms do not affect the dynamics, so we can ignore them, except for semi-classical phase diagrams (§3.4.4). To make further progress, the resulting quadratic, translation-invariant Hamiltonian must be rewritten in terms of Fourier modes: I demonstrate this process in detail for the square lattice in the next section.

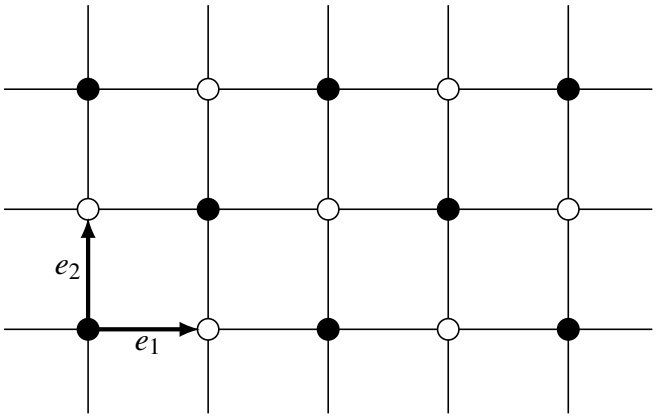

Figure 3.6. Square lattice showing the choice of basis vectors $e_{1,2}$. Figure taken from Ref. 1.

3.4 Square-lattice dimer model

Every vertex of the square lattice (Fig. 3.6) is crystallographically equivalent; its links, however, form two distinct sublattices, one each for horizontal and vertical links. Therefore, we shall label the link variables as $\delta\rho_\eta(\vec{r}+\hat{e}_\eta/2)$ and $\Phi_\eta(\vec{r}+\hat{e}_\eta/2)$, where $\eta = 1,2$ label the sublattice of horizontal and vertical links, respectively, and $\vec{r}$ is the position of a vertex; $\delta\rho$ and Φ are taken to live at the midpoints of the links. Now, introducing the Fourier decomposition

$$\delta\rho_\eta(\vec{r}+\hat{e}_\eta/2) = \frac{1}{\sqrt{N}}\sum_k \delta\rho_\eta(\vec{k})\exp\left[-i\vec{k}\cdot(\vec{r}+\hat{e}_\eta/2)\right] \tag{3.18a}$$

$$\Phi_\eta(\vec{r}+\hat{e}_\eta/2) = \frac{1}{\sqrt{N}}\sum_k \Phi_\eta(\vec{k})\exp\left[-i\vec{k}\cdot(\vec{r}+\hat{e}_\eta/2)\right] \tag{3.18b}$$

allows us to write the constraint (3.15) as

$$\sum_\eta \cos(\vec{k}\cdot\hat{e}_\eta/2)\,\delta\rho_\eta(\vec{k}) = 0. \tag{3.19}$$

It will be useful in the following to introduce the shorthand notation $c_\eta = \cos(\vec{k}\cdot\hat{e}_\eta/2)$ and $s_\eta = \sin(\vec{k}\cdot\hat{e}_\eta/2)$.

The constraint clearly imposes a relation between the two field variables $\delta\rho_1(\vec{k})$ and $\delta\rho_2(\vec{k})$. The same conclusion can also be drawn about the fields Φ_η once we notice that the Hamiltonian depends only on the specific combination of them that appears in the argument of the cosine term in (3.14):

$$\phi(\vec{R}) = \Phi_1(\vec{R}-\hat{e}_2/2) - \Phi_2(\vec{R}+\hat{e}_1/2) + \Phi_1(\vec{R}+\hat{e}_2/2) - \Phi_2(\vec{R}-\hat{e}_1/2), \tag{3.20}$$

where $\vec{R}$ is the centre of a plaquette.[†] The Fourier transform of ϕ is

$$\phi(\vec{k}) = 2\big[c_2\Phi_1(\vec{k}) - c_1\Phi_2(\vec{k})\big]. \tag{3.21}$$

Indeed, the constraints on $\delta\rho_\eta$ and on Φ_η are two sides of the same coin – conjugate variables come in pairs, so their numbers have to be the same. In our case, this is a consequence of how the RK Hamiltonian is designed: the plaquette-flipping term preserves the number of dimers at each vertex and, conversely, if one imposes a hard dimer constraint, any kinetic contribution in the Hamiltonian is projected onto a combination of loop updates, the simplest example of which is the plaquette-flipping term.

We can now rewrite the quadratic terms in (3.17) in terms of the Fourier modes (3.18) to obtain

$$\begin{aligned}\sum_k \Big\{ &- Jc_1c_2\left[\delta\rho_1(\vec{k})\delta\rho_2(-\vec{k}) + \delta\rho_2(\vec{k})\delta\rho_1(-\vec{k})\right] \\ &+ \left[(2V-J)c_2^2 + (J-V)\right]\delta\rho_1(\vec{k})\delta\rho_1(-\vec{k}) \\ &+ \left[(2V-J)c_1^2 + (J-V)\right]\delta\rho_2(\vec{k})\delta\rho_2(-\vec{k})\Big\};\end{aligned} \tag{3.22}$$

together with the Berry phase and ϕ-dependent terms, we end up with the action

$$\begin{aligned}\mathcal{S} = \int \mathrm{d}\tau \Big\{ &\sum_{k,\mu} i\,\delta\rho_\mu(\vec{k},\tau)\,\partial_\tau\Phi_\mu(-\vec{k},\tau) + \sum_{k,\mu,\nu}\frac{\mathcal{D}_{\mu\nu}(\vec{k})}{2}\delta\rho_\mu(\vec{k},\tau)\delta\rho_\nu(-\vec{k},\tau) \\ &+ \sum_R \frac{JS^2}{8}\left[1-\cos\phi(\vec{R},\tau)\right]\Big\};\end{aligned} \tag{3.23}$$

$$\mathcal{D} = 2\begin{pmatrix} J-V+(2V-J)c_2^2 & -Jc_1c_2 \\ -Jc_1c_2 & J-V+(2V-J)c_1^2 \end{pmatrix}. \tag{3.24}$$

3.4.1 Resolving the constraint

The action (3.23) is expressed in terms of the original degrees of freedom $\delta\rho$. These are, however, not the physical degrees of freedom, as their possible values are limited by the constraint (3.19). It is possible to explicitly resolve the constraint by introducing a field $h(\vec{k})$ conjugate to the physically relevant angular variable, $\phi(\vec{k})$; that is, we require

$$h(\vec{k})\partial_\tau\phi(-\vec{k}) = \sum_\eta \delta\rho_\eta(\vec{k})\partial_\tau\Phi_\eta(-\vec{k}). \tag{3.25}$$

[†] As ϕ is the lattice curl of the "vector potential" Φ (§1.1.2), it can be regarded as the magnetic field in the compact U(1) gauge theory, dual to the electric field $\delta\rho$.

Substituting (3.21) yields

$$\delta\rho_1(\vec{k}) = 2c_2 h(\vec{k}), \qquad \delta\rho_2(\vec{k}) = -2c_1 h(\vec{k}). \tag{3.26}$$

It is straightforward to verify that the field $h(\vec{k})$ automatically resolves the constraint:

$$\sum_\eta c_\eta \delta\rho_\eta = 2(c_1c_2 - c_2c_1)h(\vec{k}) = 0.$$

Once again, this result should not come as a surprise. It is a reflection, at a field-theoretic level, of the fact that the plaquette terms in the Hamiltonian respect the dimer constraint and so, if the field theory is built from plaquette kinetic terms only, the constraint is implied. In real space, the field h lives in the centres of plaquettes and is related to $\delta\rho$ as

$$\begin{aligned} \delta\rho_1(\vec{r} + \hat{e}_1/2) &= h\big[\vec{r} + (\hat{e}_1 + \hat{e}_2)/2\big] + h\big[\vec{r} + (\hat{e}_1 - \hat{e}_2)/2\big]; \\ -\delta\rho_2(\vec{r} + \hat{e}_2/2) &= h\big[\vec{r} + (\hat{e}_2 + \hat{e}_1)/2\big] + h\big[\vec{r} + (\hat{e}_2 - \hat{e}_1)/2\big]. \end{aligned} \tag{3.27}$$

This is equivalent to the standard height mapping up to a factor of z and changing the signs of h in a chequerboard pattern [1]. This shows that our construction is able to automatically recover effective representations like the height mapping.

We are now in a position to write the full large-S action for the system, including both the Berry phase and Hamiltonian contributions, in terms of the physical degrees of freedom $h(\vec{k})$ and $\phi(\vec{k})$ only. Adding more complicated ring-exchange terms to the RK Hamiltonian does not invalidate this conclusion, as the phase in each ring-exchange term can be written as a sum of ϕs over single plaquettes. Substituting the expressions of $\delta\rho_\eta(\vec{k})$ in terms of $h(\vec{k})$ and ignoring trivial constants, we obtain the action

$$\begin{aligned} \mathcal{S} = &\int \mathrm{d}\tau \sum_R \left\{ i h(\vec{R},\tau)\partial_\tau \phi(\vec{R},\tau) + \frac{JS^2}{8}\big[1 - \cos\phi(\vec{R},\tau)\big]\right\} \\ &+ \int \mathrm{d}\tau \sum_k \frac{\mathcal{D}_0(\vec{k})}{2} h(\vec{k},\tau)\, h(-\vec{k},\tau) \end{aligned} \tag{3.28}$$

$$\mathcal{D}_0 = 8\big[(2V - J)(c_1^4 + c_2^4) + (J - V)(c_1^2 + c_2^2) + 2Jc_1^2c_2^2\big]. \tag{3.29}$$

3.4.2 Photons

Ignoring for the time being the contribution to the action due to instantons between different minima of the cosine term, we can expand (3.28) about one given minimum

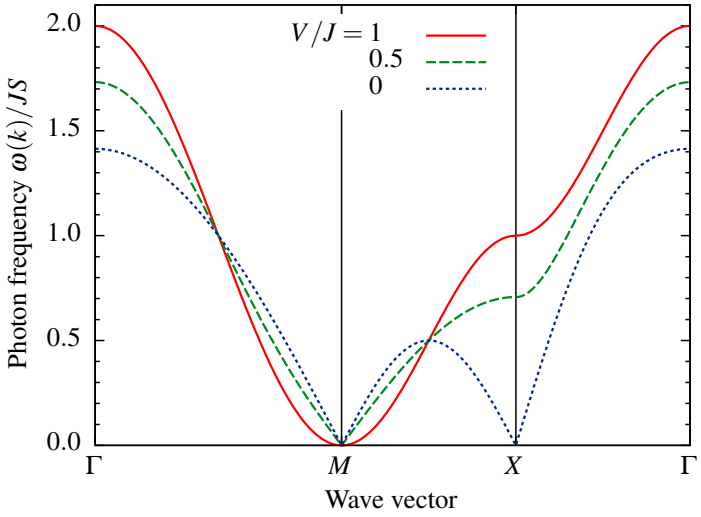

Figure 3.7. Photon dispersion relation of the large-S QDM on the square lattice. The spectrum is gapless at the $M = (\pi, \pi)$ point for all values of V/J; its dispersion is quadratic at the RK point and linear away from it. Another minimum forms at the $X = (\pi, 0)$ points for lower V, which drives an ordering transition at $V = 0$. Figure taken from Ref. 1.

(say, $\phi \equiv 0$) and integrate over ϕ to arrive at

$$\mathcal{S} = \frac{1}{2}\int \mathrm{d}\tau \sum_{k} \left[\frac{8}{JS^2}\partial_\tau h(\vec{k},\tau)\partial_\tau h(-\vec{k},\tau) + \mathcal{D}_0(\vec{k})h(\vec{k},\tau)h(-\vec{k},\tau)\right]. \tag{3.30}$$

The elementary excitations of this action are wavelike fluctuations of the field h with dispersion

$$\omega^2 = JS^2\left[(2V - J)(c_1^4 + c_2^4) + (J - V)(c_1^2 + c_2^2) + 2Jc_1^2c_2^2\right]; \tag{3.31}$$

since these fluctuations are borne out of the degrees of freedom of a U(1) gauge theory, it is natural to regard them as *emergent photons*. The dispersion (3.31) is plotted in Fig. 3.7 along high-symmetry directions in the Brillouin zone. It vanishes at (π, π) for all values of J and V. An instability develops for $V > J$, as ω^2 becomes negative near this point, leading to staggered ordering (not shown). Secondary minima appear at $(\pi, 0)$ and related points in the Brillouin zone as V is lowered: these drive the system through an instability for $V < 0$, which leads to (plaquette) dimer ordering.

A long-wavelength effective theory can be derived from (3.30) by expanding it around the gapless point (π, π):

$$\mathcal{D}_0[(\pi + q_x, \pi + q_y)] \simeq 2(J - V)(q_x^2 + q_y^2) + \frac{7V - 4J}{6}(q_x^4 + q_y^4) + Jq_x^2q_y^2$$

$$\therefore \mathcal{S} \simeq \frac{1}{2}\int \mathrm{d}\tau \mathrm{d}^2 r \left\{ \frac{8}{JS^2}(\partial_\tau \tilde{h})^2 + 2(J-V)(\nabla \tilde{h})^2 + \frac{7V-4J}{6}\tilde{h}\left(\partial_x^4 + \partial_y^4\right)\tilde{h} + J\tilde{h}\,\partial_x^2\partial_y^2\tilde{h} \right\}, \tag{3.32}$$

where $\tilde{h}$ differs from h by adding (π, π) to the wave vector: as discussed at (3.27), this recovers the standard height mapping of the square-lattice dimer model.

At the RK point, the $(\nabla\tilde{h})^2$ term vanishes, and the quartic derivative terms simplify to $J(\nabla^2\tilde{h})^2/2$, yielding the long-wavelength dispersion $\omega = k^2/(2m)$ with $m = 2/(JS)$. We note that for $S = 1$ (and $J = V = 1$), $m = 9/\pi$ [131, 146] can be obtained from the exact ground-state wave function (3.4). Our estimate, based on the lowest order of a large-S expansion, is thus within 40% accuracy; such a discrepancy at quadratic order in $1/S$ expansion is consistent, for instance, with similar calculations for quantum spin ice [17]. This can be improved by going to higher orders, and – more importantly with respect to earlier work on field theories for quantum dimer models – our approach is not limited to the fine-tuned RK point.

3.4.3 Instantons

We will now incorporate the effects of instanton events that change the field $\phi(\vec{R}, \tau)$ by 2π. These, as we shall demonstrate, always generate a mass for the photons for $V < J$, as it generally happens in 2D compact electrodynamics [16]. To this end, we are going to integrate out ϕ, taking into account the fact that the action is periodic in it.

First, we proceed by the standard Villain approach [147] and replace

$$1 - \cos\phi \quad \to \quad \frac{1}{2}\Big[\phi - 2\pi \sum_j q_j \theta(\tau - \tau_j)\Big]^2, \tag{3.33}$$

where the $q_j = q(\vec{R}_j, \tau_j)$ are integers representing instanton events, and $\theta(\tau)$ is the Heaviside step function. By integrating over ϕ and h, we obtain the action

$$\begin{aligned} \mathcal{S} &= \frac{(2\pi)^2}{2}\sum_{j,k} q(\vec{R}_j, \tau_j) G_{qq}(\vec{R}_j - \vec{R}_k; \tau_j - \tau_k) q(\vec{R}_k, \tau_k) \\ G_{qq}(\vec{k}, \omega) &= \left[\omega^2/M + (\rho_2 k^2 + \rho_4 k^4)\right]^{-1}, \end{aligned} \tag{3.34}$$

where $M = JS^2/8$ and the symbolic terms $\rho_2 k^2$ and $\rho_4 k^4$ represent the second and fourth derivative terms in the action: $\mathcal{D}_0(k) \simeq \rho_2 k^2 + \rho_4 k^4 + O(k^6)$. Equation (3.34) is the action of a Coulomb gas of charges $q = \pm 1, \pm 2, \ldots$; the fugacity of charges q is $I = \exp(-q^2 \mathcal{S}_0)$,

where $\mathcal{S}_0$ is the contribution to the action from a single instanton of charge $q = 1$:

$$\mathcal{S}_0 = \frac{1}{4\pi}\int \frac{\mathrm{d}\omega\,\mathrm{d}^2k}{\omega^2/M + (\rho_2 k^2 + \rho_4 k^4)} \approx \frac{\pi S}{8}\sqrt{\frac{J}{2\rho_4}}\,\ln(\rho_4/\rho_2) \implies I = (\rho_2/\rho_4)^{q^2\pi S/8}. \tag{3.35}$$

Since I is a rapidly decaying function of the instanton charge, we can restrict our consideration to a gas of charges $q = \pm 1$. Following Polyakov [16], we approximate the partition function of the Coulomb gas (3.34) as that of a sine-Gordon model with action (3.32), augmented by the term

$$\delta\mathcal{S} = -2\mu I\int \mathrm{d}\tau \sum_R \cos(2\pi h), \tag{3.36}$$

where $\mu\mathrm{d}\tau$ is the pre-exponential part of the instanton measure (Appendix B.2); near the RK point, $2\mu = JS^{3/2}\sqrt{\pi/2}$. The presence of this term makes the excitations massive:

$$\omega^2 = c^2k^2 + m^2, \qquad m^2 = 8\pi^2 M\mu I. \tag{3.37}$$

As we see from (3.35), this mass vanishes at the RK point.

At finite temperatures, instantons interact logarithmically:

$$E = -\frac{2\pi T^2}{\rho_2}\sum_{j<k} q_j q_k \ln\left(\frac{|\vec{R}_j - \vec{R}_k|}{r_0}\right), \tag{3.38}$$

where $r_0^2 = \rho_4/\rho_2$. The corresponding contribution to the free energy density,

$$\delta F \propto q^2 I_q \int \frac{\mathrm{d}^2 r}{r^{d_q}}, \qquad (d_q = 2\pi q^2 T/\rho_2) \tag{3.39}$$

diverges at low temperatures, where $d_1 \le 2$: above the critical temperature

$$T_c = \rho_2/\pi, \tag{3.40}$$

the cosine term in the free energy functional of static fluctuations,

$$F = \int \mathrm{d}^2 r\,\left[\frac{\rho_2}{2}(\nabla h)^2 + \frac{\rho_4}{2}(\nabla^2 h)^2 - 2\mu I\cos(2\pi h)\right], \tag{3.41}$$

is irrelevant. Above T_c, therefore, we have a critical phase; below it, the correlation

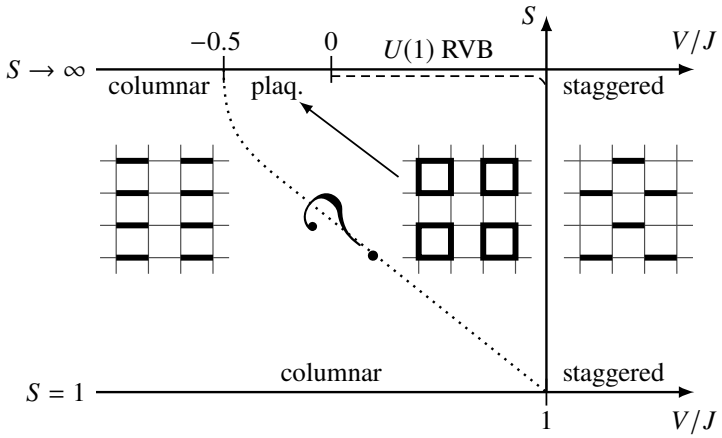

Figure 3.8. Phase diagram of the square lattice quantum dimer model at $S = \infty$ (this work) and $S = 1$ [138, 149–156]. Figure taken from Ref. 1.

length of the h field is finite:

$$\xi \sim r_0 (\mu I / T_c)^{-1/(2-d_1)}. \tag{3.42}$$

The divergence of ξ at T_c corresponds to the melting of the instanton-generated valence bond crystal via a Berezinskii–Kosterlitz–Thouless transition [148].

3.4.4 Large-S phase diagram

In the $S \to \infty$ limit, zero-point fluctuations of any soft modes are negligible, so ρ and Φ can be treated as commuting, classical variables: the ground state energy of the system is therefore given by classical minimisation of the Hamiltonian (3.14). Since $\sqrt{\rho\rho\rho\rho} \geq 0$ always, the Φ in such an optimal state satisfy $\cos(\Phi_1 - \Phi_2 + \Phi_3 - \Phi_4) = 1$ for all plaquettes, which is achieved, for instance, by setting $\Phi \equiv 0$.

Finding the optimal values of ρ in full generality is more difficult. However, one can always compare the ground state energies of phases suggested in the literature, or develop a variational ansatz that captures several such phases. In the case of the square lattice, we considered states in which ρ is constant within each set of bonds populated in the four columnar ordered states. Such an ansatz can describe columnar and plaquette ordered states as well as the RVB liquid.

Comparing the ground-state energies of these phases and staggered order yields the $S \to \infty$ phase diagram shown in Fig. 3.8. As expected, the ground state is staggered for $V > J$ and columnar at $V \to -\infty$. At intermediate V/J, we see a plaquette ordered phase as well as an extended U(1) RVB liquid, with phase boundaries corresponding to the instabilities shown in Fig. 3.7. The latter is unstable at finite S due to instanton

effects, as discussed above. The fate of the plaquette phase is less clear: since it is ordered, instantons are unlikely to substantially affect its stability, so the evolution of the columnar–plaquette phase boundary is determined by lattice effects. It may well be possible that the plaquette order survives at $S > 1$ and has a proximity effect near the RK point even at $S = 1$. This could explain why numerical simulations of the square-lattice dimer model struggle to establish its true ground state in this regime [138, 149–156].

3.5 Other lattices

Beyond the square-lattice dimer model, our approach can straightforwardly be applied to RK Hamiltonians on other bipartite lattices, both in two and three dimensions. In this section, I demonstrate this on QDMs on the cubic, honeycomb, and diamond lattices. Finally, I discuss the applicability of the semiclassical approach to $\mathbb{Z}_2$ liquids on the non-bipartite triangular lattice.

3.5.1 Cubic lattice

The RK Hamiltonian on the cubic lattice can be thought of as the sum of square-lattice dimer models (3.9) in the three sets of $\{100\}$ crystallographic planes, linked by the shared large-S constraint (3.12). Similar to the square-lattice case, we introduce the radial gauge (3.13), setting $z = 6$, and expand the resulting Hamiltonian (3.14) in $\delta\rho$. After introducing the Fourier modes (3.18), the large-S constraint can be written as

$$\sum_{\eta} c_\eta \delta\rho_\eta(\vec{k}) = 0, \qquad [c_\eta = \cos(\vec{k}\cdot\hat{e}_\eta)/2)] \tag{3.43}$$

while the action becomes

$$\mathcal{S} = \int \mathrm{d}\tau \Big\{ \sum_{k,\mu} i\,\delta\rho_\mu(\vec{k},\tau)\,\partial_\tau \Phi_\mu(-\vec{k},\tau) + \sum_{k,\mu,\nu} \frac{\mathcal{D}_{\mu\nu}(\vec{k})}{2} \delta\rho_\mu(\vec{k},\tau)\delta\rho_\nu(-\vec{k},\tau) + \sum_{R,\mu} \frac{JS^2}{18}\big[1-\cos\phi_\mu(\vec{R},\tau)\big]\Big\} \tag{3.44}$$

to quadratic order in $\delta\rho$, where

$$\mathcal{D} = 2\begin{pmatrix} (2V-J)(c_2^2+c_3^2) & -Jc_1c_2 & -Jc_1c_3 \\ -Jc_1c_2 & (2V-J)(c_1^2+c_3^2) & -Jc_2c_3 \\ -Jc_1c_3 & -Jc_2c_3 & (2V-J)(c_1^2+c_2^2) \end{pmatrix} + 4(V-J)\mathbb{1}, \tag{3.45}$$

and ϕ_μ labels the arguments of the cosine terms for plaquettes normal to the $\hat{e}_\mu$ direction: unlike the case of the square lattice, these sets of plaquettes are crystallographically inequivalent.

We could now rewrite the action in terms of the physical degrees of freedom only, that is, replace the three $\delta\rho$ fields with two "height fields" that impose the constraint (3.43) automatically.[†] However, doing so is much more algebraically tedious than for the square lattice, as there is no geometrically well-motivated choice of the fields. In fact, the $\delta\rho$ and ϕ are electric and magnetic fields in a lattice U(1) gauge theory (§1.1.2, §2.1.1), which are most naturally expressed in terms of a redundant vector field obeying divergence-free constraints.

Photons

Therefore, we derive the dispersion of photons in the U(1) gauge theory using the action (3.44) in terms of the $\delta\rho$ and Φ fields. In the first approximation, we again neglect instantons in the fields ϕ and expand them to quadratic order around the minimum $\phi \equiv 0$:

$$\frac{JS^2}{18}\sum_{R,\mu}\Big[1-\cos\phi_\mu(\vec{R},\tau)\Big] \simeq \frac{JS^2}{36}\sum_{R,\mu}\phi_\mu(\vec{R},\tau)^2, \tag{3.46}$$

which allows us to write the last term of (3.44) in terms of Fourier modes. Similar to (3.21), we have $\phi_\mu(\vec{k}) = 2\varepsilon_{\mu\nu\lambda}c_\lambda\Phi_\nu(\vec{k}) \equiv \mathcal{Z}_{\mu\nu}\Phi_\nu(\vec{k})$. $\mathcal{Z}$ is a nonzero, traceless, antisymmetric matrix; it has a zero eigenvalue, $\mathcal{Z}_{\mu\nu}c_\nu = 0$, which will be useful to project out the unphysical mode that does not obey the constraint (3.43). The action (3.46) can now be written as

$$\frac{JS^2}{36}\sum_\mu \phi_\mu(\vec{k})\phi_\mu(-\vec{k}) = -\frac{JS^2}{36}\sum_{\mu\nu}[\mathcal{Z}^2(\vec{k})]_{\mu\nu}\Phi_\mu(\vec{k})\Phi_\nu(-\vec{k}) \equiv \sum_{\mu\nu}\frac{\mathcal{M}_{\mu\nu}}{2}\Phi_\mu(\vec{k})\Phi_\nu(-\vec{k}); \tag{3.47}$$

after integrating the fields Φ out from (3.44), we get

$$\mathcal{S} = \frac{1}{2}\int \mathrm{d}\tau \sum_{k,\mu,\nu}\Big[\mathcal{M}^{-1}_{\mu\nu}\,\partial_\tau\delta\rho_\mu(\vec{k})\,\partial_\tau\delta\rho_\nu(-\vec{k}) + \mathcal{D}_{\mu\nu}\,\delta\rho_\mu(\vec{k})\,\delta\rho_\nu(-\vec{k})\Big]. \tag{3.48}$$

The photon dispersion is given by the eigenvalues of $\mathcal{M}\mathcal{D} = -2JS^2\mathcal{Z}^2\mathcal{D}/9$, after projecting out the unphysical modes that do not satisfy the constraint $\sum_\mu c_\mu\delta\rho_\mu = 0$. Formally, this could be done by adding an infinite Lagrange multiplier, but there is no need to do

[†]The corresponding reduction in degrees of freedom for the ϕ fields comes from the fact that, as lattice curls of Φ, their lattice divergence must be zero for each cube of the lattice. In Fourier space, this implies $\sum_\mu c_\mu\phi_\mu(\vec{k}) = 0$.

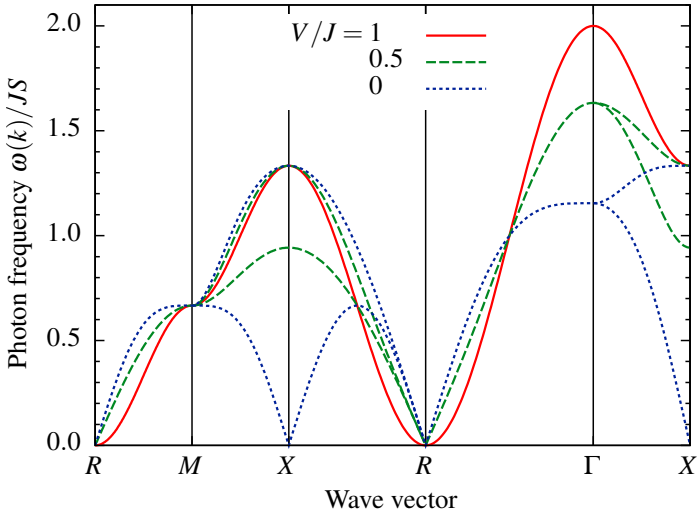

Figure 3.9. Photon dispersion relation of the large-S QDM on the cubic lattice. The spectrum is gapless at the $R = (\pi, \pi, \pi)$ point; its dispersion is quadratic at the RK point and linear away from it. The spectrum has two non-degenerate branches away from the RK point. The lower branch develops minima at the $X = (\pi, 0, 0)$ points for lower values of V, which drives an ordering transition at $V = 0$. Figure taken from Ref. 1.

so because the only unphysical mode is a zero mode of $\mathcal{Z}$, and hence of $\mathcal{Z}^2\mathcal{D}$.[†] The two nonvanishing eigenvalues are

$$\omega^2 = \frac{4JS^2}{9}\Big\{J\left[2(c_1^2+c_2^2+c_3^2) - (c_1^4+c_2^4+c_3^4)\right] + 2V\left[c_1^4+c_2^4+c_3^4 - (c_1^2+c_2^2+c_3^2) + c_1^2c_2^2 + c_1^2c_3^2 + c_2^2c_3^2\right] \pm 2|J-V|\sqrt{c_1^4c_2^4 + c_1^4c_3^4 + c_2^4c_3^4 - c_1^2c_2^2c_3^2(c_1^2+c_2^2+c_3^2)}\Big\}, \tag{3.49}$$

which are plotted for three values of V/J in Fig. 3.9. It is interesting to note that $\omega^2|_{J=V} \propto (c_1^2 + c_2^2 + c_3^2)^2$: the two bands are degenerate at the RK point, with quadratic dispersion around the gapless point (π, π, π). Expanding near this point for $V \neq J$, we find

$$\omega^2 \simeq \frac{2S^2}{9} J(J-V)k^2, \tag{3.50}$$

where $\vec{k}$ is the (small) vector distance from (π, π, π): the speed of light is $c = \sqrt{2J(J-V)}\, S/3$.

[†]In fact, the matrix $\mathcal{M}^{-1}$ in (3.48) diverges because of the same zero mode. Formally, a small $\lambda c_\mu c_\nu$ component should be added to $\mathcal{M}$ to make it invertible, and the unphysical mode would be projected out by taking $\lambda \to 0$. However, as the photon spectrum is given by that of the well-behaved matrix $\mathcal{M}\mathcal{D}$, we achieve the same results by simply dropping its zero eigenvalue.

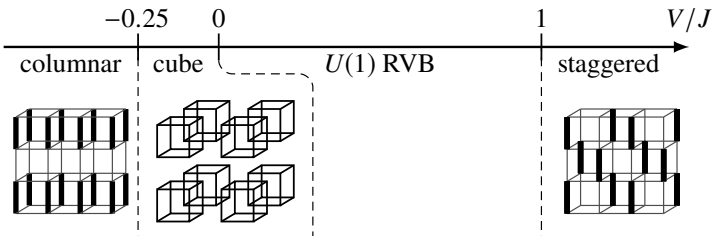

Figure 3.10. $S \to \infty$ phase diagram of the cubic lattice quantum dimer model. To my knowledge, there are no conclusive studies of the $S = 1$ QDM on the cubic lattice, to which one could compare the large-S results of this work. Figure taken from Ref. 1.

Instantons

The instanton contributions are calculated similarly to the square lattice case, by performing the Villain transformation (3.33) on ϕ_μ in (3.44) and integrating out the smooth fields $\Phi_\mu, \delta\rho_\mu$. The result is a 3D Coulomb gas action for the integer charges q_μ,

$$\mathcal{S} = \frac{(2\pi)^2}{2} \sum_{\omega,k} \sum_{\mu,\nu} q_\mu(-\omega,-k) \left[\frac{18\omega^2}{JS^2} \mathcal{M}^{-1} + \mathcal{D} \right]^{-1}_{\mu\nu} q_\nu(\omega,k), \tag{3.51}$$

with the standard unscreened long-range Coulomb interaction. The constraint (3.43) on the fields ϕ_μ, however, implies the equivalent constraint

$$c_1 q_1 + c_2 q_2 + c_3 q_3 = 0 \tag{3.52}$$

on the instanton configurations $q_\mu(\omega,k)$ allowed in the low-energy sector.[†] In 3 + 1 dimensions, such instantons are irrelevant [16] and can safely be neglected.

Large-S phase diagram

Using the method described in §3.4.4, we can obtain the ground state of the QDM Hamiltonian (3.9) in the limit $S \to \infty$ (Fig. 3.10). Similar to the square-lattice case, we observe an extended U(1) RVB liquid phase: in three dimensions, this phase is expected to survive at $S = 1$. For $V > J$, the photon modes become unstable at the (π,π,π) point, leading to staggered order. Likewise, the instability of the $(\pi,0,0)$ points for $V < 0$ drives a transition into an RVB solid phase with isolated, resonating cubes. For $V < -J/4$, this phase gives way to columnar order.

[†]Instantons that do not obey this constraint also exist and give rise to pairs of Dirac-quantised magnetic monopoles, analogous to visons in quantum spin ice [15]. These monopoles are gapped quasiparicles and only renormalise the low-energy theory through ring-exchange processes that obey the constraint (3.52).

3.5.2 Honeycomb and diamond lattices

The calculations for these lattices are straightforward generalisations of the previous two cases with some minor differences:

- The smallest plaquettes on which ring-exchange terms can be defined consist of six links: therefore, the Hamiltonian (3.9) is replaced by

$$H_D = \sum_{\hexagon} \left[-J \left(\hat{b}_1^\dagger \hat{b}_3^\dagger \hat{b}_5^\dagger \hat{b}_2 \hat{b}_4 \hat{b}_6 + \text{H.c.} \right) + V \left(\hat{b}_1^\dagger \hat{b}_1 \hat{b}_3^\dagger \hat{b}_3 \hat{b}_5^\dagger \hat{b}_5 + \hat{b}_2^\dagger \hat{b}_2 \hat{b}_4^\dagger \hat{b}_4 \hat{b}_6^\dagger \hat{b}_6 \right) \right]. \tag{3.53}$$

  In the radial gauge (3.13), (3.53) can still be expanded to quadratic order in $\delta\rho$ much the same way as for the square lattice.

- There are two inequivalent lattice sites in each unit cell (marked black and white in Fig. 3.11) that give rise rise to two distinct large-S constraints (3.12). In the honeycomb lattice, there are three links per unit cell, leaving one physical degree of freedom, which allows us to define a height mapping analogous to (3.26). In the diamond lattice, there are four links, and hence two physical degrees of freedom, per unit cell: this and the complicated structure of the lattice requires us to resolve the constraint implicitly (§3.5.1).

The algebra required to obtain the photon spectrum in both cases is rather tedious and is presented in Ref. 1 in full detail. The spectra are shown along high-symmetry direction for several values of V/J in Fig. 3.11. They are both gapless at the Γ point, with the expected quadratic and linear dispersions at and away from the RK point, respectively. In particular, the speed of light in the diamond lattice QDM is $c = S^2\sqrt{J(J-V)/6}$, compared to the quantum Monte Carlo result $c \approx \sqrt{0.8J(J-V)}$ for $S = 1$ [158, 159]. We expect higher-order corrections in the $1/S$ expansion to substantially reduce this discrepancy [17].

Phase diagrams of the $S \to \infty$ models were obtained as a function of V/J by comparing the classical ground-state energies of several ordered and resonating phases suggested in the literature (Refs. [157, 160] for the honeycomb lattice, Refs. [19, 158, 159] for the diamond lattice). The results are shown in Fig. 3.11, together with the corresponding $S = 1$ phase diagrams [157, 158]. In both cases, the liquid phase remains stable beyond the transition point to a plaquette or columnar ordered phase, indicating a first-order transition. For the diamond lattice, no resonating solid phase is identified in either limit: both phase diagrams consist of staggered ordered, U(1) liquid and ordered R [159] phases. The U(1) RVB liquid phase is destroyed by instantons on the 2D honeycomb lattice for finite S (§3.4.3), but survives on the 3D diamond lattice. Nevertheless, it becomes

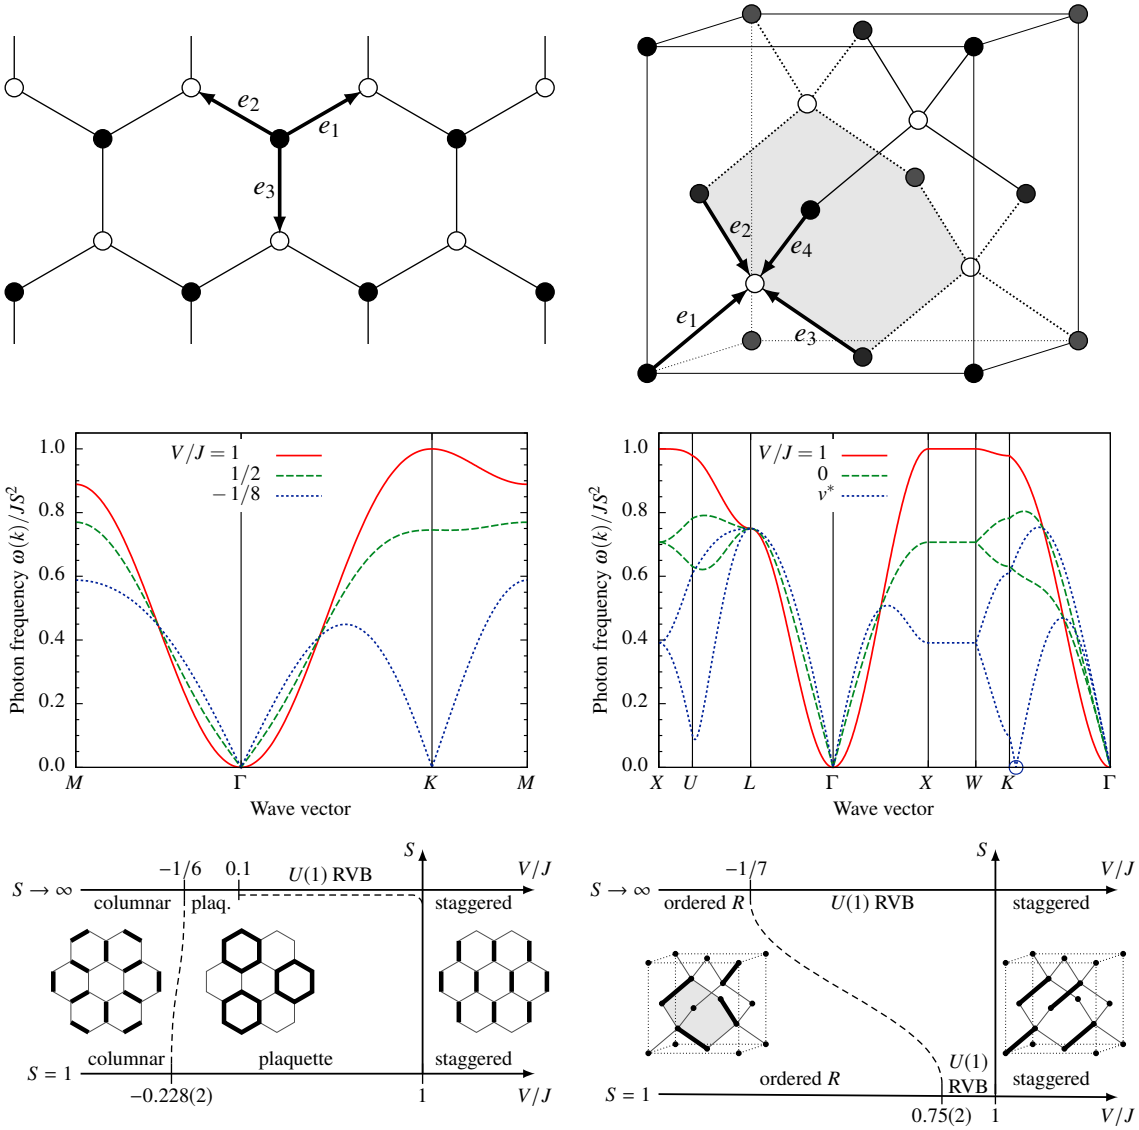

Figure 3.11. **Left**: Large-S photon dispersion relation and phase diagram of the honeycomb-lattice QDM. The spectrum is gapless at the Γ point in the Brillouin zone. Another minimum develops for small V at the K points; however, a first-order plaquette ordering transition occurs before this minimum would become unstable. The $S = 1$ phase diagram is taken from Ref. 157.
Top right: Unit cell of the diamond lattice. The two interpenetrating FCC sublattices are shown as solid and open dots, respectively. A hexagonal plaquette is shown shaded by way of example.
Right: Large-S photon dispersion relation and phase diagram of the diamond-lattice QDM. The spectrum is gapless at the Γ point in the Brillouin zone and has two non-degenerate branches away from the RK point. Another minimum develops for small V near the K points; however, a first-order columnar ordering transition occurs before this minimum would become unstable at $V/J = v^* \approx -0.694$. The $S = 1$ phase diagram is taken from Ref. 158.
Figures taken from Ref. 1.

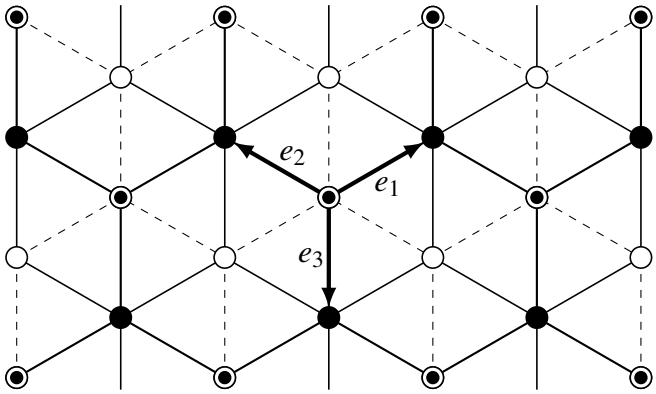

Figure 3.12. Projecting a cubic lattice down a ⟨111⟩ axis onto a triangular lattice. The figure shows a single layer of cubes: two vertices in each cube are projected onto one another (circled black dots), which originally belonged to different sublattices of the bipartite cubic lattice (black and white dots). The ⟨100⟩ unit vectors of the cubic lattice are projected onto the linearly dependent lattice vectors $\hat{e}_{1,2,3}$ of the triangular lattice. Figure taken from Ref. 1.

far narrower at $S = 1$ than at $S \to \infty$, consistent with the intuition that soft dimers (and spins) favour fluctuating phases (§2.2).

3.5.3 Outlook: non-bipartite lattices

Our method has only been applied to bipartite lattices in the discussion above. This is a natural choice, since the dimer density operator ρ and the conjugate angular variable Φ introduced in the radial gauge (3.13) can straightforwardly be interpreted as a compact U(1) electric field and magnetic vector potential, respectively [15]; furthermore, the focus of $\mathbb{Z}_2$ gauge theories on parities, rather than the dimer numbers themselves, stands at odds with our approach that is underpinned by coherent states with a large dimer number variance. Nevertheless, the procedure outlined in §3.3 can be performed on non-bipartite lattices, leading to large-S descriptions of RK models on the same.

A particularly interesting case is that of the triangular lattice, which can be derived from the cubic lattice by projecting it down a ⟨111⟩ axis (Fig. 3.12). The triangular-lattice RK model can be derived from the cubic one using the same projection, so long as all links projected onto one another are in the same bosonic coherent state b. The derivation of the cubic-lattice dispersion relation (3.49) carries over unchanged, after identifying the components $k_{1,2,3}$ of the wave vector with the linearly dependent $\vec{k} \cdot \hat{e}_{1,2,3}$: the triangular-lattice dispersion relation is a 2D cut of the cubic one in the plane $k_1 + k_2 + k_3 = 0$. Since the gapless point of the latter, (π, π, π), is not in this plane, the resulting "photon" dispersion relation is gapped from $V/J = -1/2$ to above 1 [Fig. 3.13(a)]; nevertheless, we still expect a first-order transition into a staggered ordered phase at the RK point.

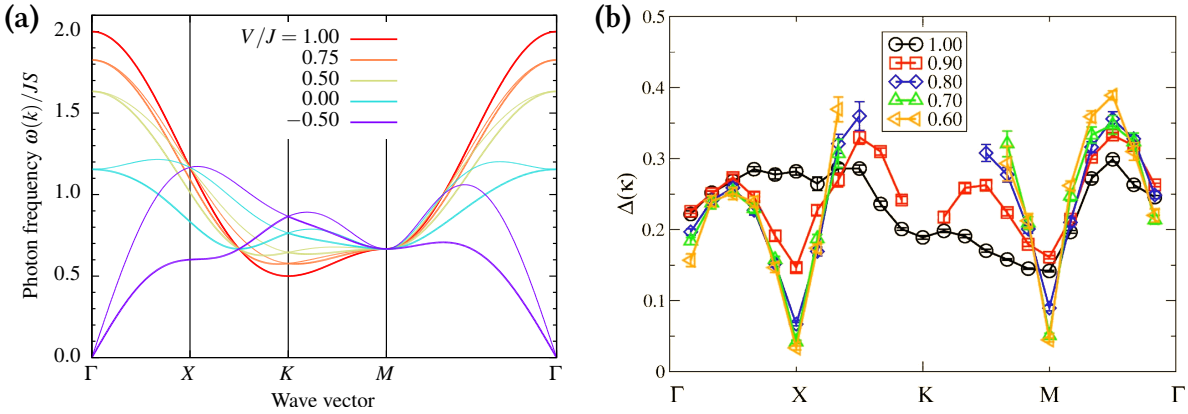

Figure 3.13. **(a)** Semiclassical "photon" dispersion relation in the $S \to \infty$ triangular-lattice RK model. It becomes gapless at the Γ point in the Brillouin zone at $V/J = -1/2$, but remains gapped above that up to the RK point.
(b) Energy of the lowest "non-vison-like" [137] excitation in the $S = 1$ triangular-lattice RK model, obtained from Green's function Monte Carlo on a 432-site cluster [143]. It is also gapped at the RK point, but develops soft modes at the X and M points around $V/J = 0.8$. Figure taken from Ref. 143.

Visons can also be captured using the semiclassical description. In the path-integral formalism, the vison creation operator (3.7) turns into

$$(-1)^{\hat{N}_{AB}} = \exp\left(i\pi \sum_{i=A}^{B} \hat{n}_i\right) \longrightarrow \exp\left(i\pi \sum_{i=A}^{B} \rho_i\right) = \prod_{i=A}^{B} e^{i\pi\rho_i}; \tag{3.54}$$

since the angular variable Φ_i is canonically conjugate to ρ_i, the effect of this operator is to change Φ along a path between plaquettes A and B by π.[†] As a result, the circulation of Φ[‡] along any loop that encloses a single vison changes by π, recovering the semionic mutual statistics with monomers (§3.2.1). The phenomenology of these semiclassical visons is similar to those in large-S quantum spin ice (§2.3.3): their energy is on the order of Hamiltonian terms, $O(JS^2)$, parametrically larger than that of the photons, $O(JS)$; therefore, in a semiclassical numerical simulation, photons would appear as normal modes, while visons remain quasiparticles with a finite energy.

This situation is qualitatively different from the $S = 1$ QDM, where visons are the lowest-energy excitations [137], whose condensation around $V/J = 0.8$ drives the transition into a resonating-plaquette phase called $\sqrt{12}\times\sqrt{12}$ order [134, 142, 143, 161]. Indeed, soft modes for $S = 1$ are seen to form at the X and M points of the Brillouin zone [143]

[†]The same conclusion is also reached by observing that the coherent states $|\beta\rangle$ and $|-\beta\rangle$ differ precisely by the sign of odd parity Fock states, cf. (3.10).

[‡]That is, the sum of Φ along an even-length loop with alternating signs. As the lattice is non-bipartite, a given Φ may appear with different signs in different loops. However, as $-\pi$ and π are the same phase, the π-shift that defines the vison is independent of such choices.

[Fig. 3.13(b)], which are consistent with $\sqrt{12} \times \sqrt{12}$ order; by contrast, the instability at the Γ point in semiclassics appears to herald a columnar order. Since visons are large perturbations of the semiclassical RVB liquid state around which the expansion (3.13) is performed, it is difficult to tell whether precursors to their $S = 1$ behaviour exist in the $S \to \infty$ limit: this question may be studied more easily using numerical simulations along the lines of §2.3. Even so, the dynamics of large-S visons would be due to instanton effects in the Φ field [97]: these are not captured by the same simulations, and are generally suppressed exponentially in S. These probably put a detailed reconstruction of the $S = 1$ physics out of the reach of the semiclassical approach.

Nevertheless, the construction of "photons" on the triangular lattice can be used to propose variational "photon" states in the $S = 1$ RK model. Comparing the properties of these states to the variational vison states (3.7) could shed light on the relationship between U(1) and $\mathbb{Z}_2$ RVB liquids and, particularly, on the mechanism by which the U(1) gauge theory breaks down into a $\mathbb{Z}_2$ one [162]. Owing to the simple structure of ground states at the RK point, these studies may be tractable analytically, or done exactly using classical Monte-Carlo techniques [60, 135], offering, in principle, more rigorous results than usually possible for spin-liquid systems.

4

Neural-network wave functions and the sign problem

Over the last decade, machine learning has had a profound impact on nearly all aspects of life as well as the physical sciences [163], mostly owing to the ability of neural networks to discover patterns in complex, data with very high dimensionality [164]. The same properties mark neural networks out as general-purpose wave-function ansätze, with potential to overcome the limitations of existing approaches, such as (path-integral) quantum Monte Carlo or tensor networks. In this chapter, I discuss neural quantum states (NQS), their relationship with the Monte-Carlo sign problem, and our work [3] on solving the NQS sign problem for unfrustrated antiferromagnets and alleviating it for frustrated ones.

4.1 Introduction to quantum Monte Carlo

God does not play dice with the universe.

— Albert Einstein

The central problem of simulating quantum mechanics numerically is dealing with the enormous amount of information contained in the quantum wave function. For instance, simulating the time evolution of a system of N classical particles requires solving $O(N)$ coupled ordinary differential equations (ODEs) (cf. §2.2); by contrast, the same task for a system of N spin-1/2 degrees of freedom leads to 2^N ODEs, while the energy levels of the same system follow from diagonalising a $2^N \times 2^N$ matrix (from here on, we focus on this latter task). The exponentially growing computational complexity means that not even moderately large quantum systems can be simulated exactly on a classical computer [165]. Accessing larger system sizes requires approximations along the following strategies:

- Wave functions are represented using an ansatz whose observable expectation values can be evaluated exactly with modest computational effort. The variational energy of such a parametrised wave function can be minimised explicitly, leading quickly to the best possible approximation of the ground state with that ansatz. The most successful examples of this approach are *tensor networks* [166], for which variational energies and other local observables can be evaluated efficiently by repeated contractions of small (relative to the Hilbert space dimension) tensors. Being variational, the accuracy of these methods is limited by the ability of the ansatz to represent physically relevant wave functions; indeed, representing an arbitrary wave function requires exponentially many parameters, leading to no improvement over exact diagonalisation.

- A subset of the Hilbert space is used to estimate operator expectation values. Since this subset is selected by importance-sampling Monte-Carlo algorithms, these approaches are generally called *quantum Monte Carlo (QMC)*. A great multitude of QMC techniques has been developed [38], differing in what is sampled by the Monte-Carlo algorithm, as well as in matters of implementation. Here, I review *variational Monte Carlo,* which forms the basis of our work with neural-network wave functions, and *path-integral Monte Carlo* to demonstrate the origins of the *Monte-Carlo sign problem.*

4.1.1 Variational Monte Carlo (VMC)

No great discovery was ever made without a bold guess.
— Sir Isaac Newton

Consider a Hilbert space with a complete, orthonormal basis (the *computational basis*). For spin systems, the typical computational basis consists of the joint eigenstates of all σ^z operators; for $S = 1/2$, these can be labelled with a bit string of + and − signs. The wave function is now a function that maps these bit strings to complex numbers; a *wave-function ansatz* is a family of such functions parametrised by some other parameters θ:

$$|\Psi_\theta\rangle = \sum_{\boldsymbol{\sigma}} \Psi(\boldsymbol{\sigma};\theta)|\boldsymbol{\sigma}\rangle, \tag{4.1}$$

where $\boldsymbol{\sigma}$ runs over the σ^z basis states. Importantly, the number of parameters is much smaller than the Hilbert space dimension, so not all possible wave functions can be represented exactly: the accuracy of the converged ground state depends strongly on the variational ansatz.

Variational Monte-Carlo approaches aim to approximate the ground state of a Hamiltonian H by an appropriate choice of the parameters θ. As a proxy for the true ground state, we minimise the variational energy

$$E(\theta) = \frac{\langle \Psi_\theta | H | \Psi_\theta \rangle}{\langle \Psi_\theta | \Psi_\theta \rangle}, \tag{4.2}$$

which has a global minimum at the ground state. Since we only assume basic continuity properties about $\Psi(\boldsymbol{\sigma}; \theta)$, this minimisation can only be done gradually.[†] There are two common strategies for doing so:

Stochastic gradient descent (SGD)

Monte-Carlo sampling is used to estimate the derivatives of the energy (4.2) with respect to the parameters θ and the parameters are moved along the direction of steepest descent: $\theta_i \mapsto \theta_i - \eta \partial_{\theta_i} E(\theta)$, where η is the *learning rate*. This approach is popular for machine-learning tasks because it can handle millions of parameters efficiently; however, it is ill-suited to optimising variational wave functions [167, 168].

Stochastic reconfiguration (SR)

A strategy that greatly improves the convergence of variational wave functions compared to SGD is approximating their imaginary-time evolution rather than directly minimising the variational energy [38, 168]. Namely, given a set of parameters θ, we want to find $\theta' = \theta + \delta\theta$ such that $|\Psi_{\theta'}\rangle$ have a high overlap with $e^{-\eta H}|\Psi_\theta\rangle \approx (1 - \eta H)|\Psi_\theta\rangle$ for a sufficiently small η. The ground state has a higher overlap with the latter than with $|\Psi_\theta\rangle$, so iterating this procedure will converge to the ground state.

To first order in η, the optimal $\delta\theta$ is given by (Appendix C.1)

$$\sum_j \underbrace{\operatorname{Re}\operatorname{cov}(O_j, O_k)}_{S_{kj}} \delta\theta_j = -\eta \underbrace{\operatorname{Re}\operatorname{cov}(E_{\mathrm{loc}}, O_k)}_{\partial_{\theta_k} E(\theta)}. \tag{4.3}$$

Here, $O_j(\boldsymbol{\sigma}) = \left\langle \boldsymbol{\sigma} \middle| \partial_{\theta_j} \Psi \right\rangle / \langle \boldsymbol{\sigma} | \Psi \rangle = \partial_{\theta_j} \log \langle \boldsymbol{\sigma} | \Psi \rangle$; $E_{\mathrm{loc}}(\boldsymbol{\sigma}) = \langle \boldsymbol{\sigma} | H | \Psi \rangle / \langle \boldsymbol{\sigma} | \Psi \rangle$ is a local estimate of the variational energy in the sense that

- its average weighted according to the quantum probability distribution $p(\boldsymbol{\sigma}) = \left|\langle \boldsymbol{\sigma} | \Psi \rangle\right|^2 / \langle \Psi | \Psi \rangle$ is the variational energy (4.2);
- in an exact eigenstate, it is equal to the corresponding energy eigenvalue for all $\boldsymbol{\sigma}$.

[†]For certain ansätze (e.g., tensor networks), very efficient algorithms can be designed by exploiting their structure, which bypass Monte-Carlo sampling altogether.

The covariances in (4.3) are evaluated with respect to the quantum probability distribution $p(\boldsymbol{\sigma})$, so they can readily be estimated with importance-sampling Monte-Carlo techniques, such as the Metropolis–Hastings algorithm.

Equation (4.3) differs from the gradient-descent prescription, $\delta\theta = -\eta\partial_{\theta_k}E(\theta)$, by introducing the covariance matrix S. This matrix depends entirely on the parametrisation of the wave function rather than its energy under the Hamiltonian, so it can be thought of as a metric tensor on the parametrised Hilbert space [38]. Therefore, stochastic reconfiguration is analogous to the natural gradient approaches used to stabilise gradient descent in other machine learning contexts [167].

Common ansätze for variational Monte Carlo

As discussed already, the success of VMC (and any variational principle) depends crucially on the ability of the chosen variational ansatz to capture physically relevant quantum states. In this section, I review some conventional ways of constructing ansätze for electronic and spin systems on lattices.

The most common starting point for systems of interacting electrons (e.g., Hubbard models) is a *mean-field wave function,* that is, the ground state of a non-interacting Hamiltonian, composed of tight-binding hopping and pairing[†] terms, the values of which are treated as variational parameters. For a regular lattice, the ground state can be written down explicitly in reciprocal space by solving single-body problems with small Hilbert spaces at each wave vector. Interactions can be accounted for more easily in the occupation-number basis in real space, where the mean-field wave function is given by a single Slater determinant, analogous to continuum Hartree–Fock theory. Evaluating this Slater determinant directly is computationally expensive; however, derivatives with respect to the entries [used to evaluate $O(\boldsymbol{\sigma})$] and ratios of mostly identical Slater determinants [used in the Metropolis condition and to evaluate $E_{\mathrm{loc}}(\boldsymbol{\sigma})$] can be computed more efficiently by exploiting linear-algebra identities, making VMC with mean-field ansätze viable [38].

We can use the same mean-field wave functions to model antiferromagnets[‡] by observing that a half-filled Hubbard model in the limit of infinitely strong repulsive interactions reduces to the Heisenberg model [32] (see §1.2 for a more detailed account of this connection). The infinitely strong interactions force each site to host precisely

[†]Pairing terms break the fermion-number conservation of realistic interacting Hamiltonians. This can be rectified by projecting the mean-field wave function onto basis states with a fixed electron number: in VMC, this amounts to sampling those states only.

[‡]Ferromagnetic Hamiltonians have no sign problem, so they are usually studied using path-integral Monte Carlo and other unbiased approaches.

one particle, which cannot be achieved directly with mean-field wave functions: this can be corrected for by projecting out all basis states with unoccupied or doubly occupied sites; this *Gutzwiller projection* (Fig. 1.4) can easily be implemented in VMC by not sampling these basis states. Gutzwiller-projected mean-field states are very well suited to quantum-spin-liquid states as the fermionic hopping and pairing terms (1.20) do not break SU(2) spin-rotation symmetry; ordered states can be modelled by coupling the mean field Hamiltonian to an external magnetic field, used as a variational parameter.[†]

Converged mean-field wave functions provide significant insight into the physics of the ground state: a nonzero pairing term might indicate superconductivity or $\mathbb{Z}_2$ topological order [7, 36, 169, 170]; ordering transitions may be detected through symmetry-breaking mean-field terms [171]; and for spin systems, ansätze with different spin quantum numbers can be constructed, giving direct access to spin gaps [172–175]. On the other hand, electronic mean-field states do not capture electron–electron correlations and thus result in poor approximations of the ground state. The Gutzwiller projection yields nontrivial correlations, making them more applicable to spin liquids; however, they are still not flexible enough to represent the true ground states.

The performance of mean-field ansätze can systematically be improved using *many-body correlators:* these are operators that are diagonal in the computational basis, so they can be applied to a variational wave function by multiplying wave-function amplitudes with the corresponding matrix elements. For spin systems, the simplest choice is the *Jastrow wave function*

$$|\text{Jastrow}\rangle = \exp\Big(\sum_{ij} J_{ij}\hat{\sigma}_i^z\hat{\sigma}_j^z\Big)|\text{MF}\rangle = \sum_{\boldsymbol{\sigma}} |\boldsymbol{\sigma}\rangle\langle\boldsymbol{\sigma}|\text{MF}\rangle \exp\Big(\sum_{ij} J_{ij}\sigma_i^z\sigma_j^z\Big), \tag{4.4}$$

where the J_{ij} are variational parameters. These wave functions capture many-body correlation effects more accurately and provide significantly better variational energies.

One can also construct a wave-function ansatz directly in terms of the computational basis states, without reference to a mean-field wave function. A simple but remarkably successful approach is due to Huse and Elser [176], who found that wave functions of the form

$$|\text{HE}\rangle = \sum_{\boldsymbol{\sigma}} \exp\Big(\sum_{ij} J_{ij}\sigma_i^z\sigma_j^z + \sum_i h_i\sigma_i^z\Big)|\boldsymbol{\sigma}\rangle, \tag{4.5}$$

attain low variational energies for Heisenberg antiferromagnets on both the square and the triangular lattice. Here, the J_{ij} and h_i are complex variational parameters. For $J = 0$, the ansatz describes product states without any entanglement; the bilinear terms intro-

[†]The coupling term to an external field, $\mathbf{h}\cdot\mathbf{S}$, is a fermion bilinear, so the mean-field ground state remains a Slater determinant.

duce long-range two-spin correlations in a multiplicatively extensive way. Importantly, Eq. (4.5) can naturally represent the Marshall sign rule [177] for unfrustrated antiferromagnets exactly (§4.1.4). The Huse–Elser construction can naturally be extended to allow for correlations between clusters of more than two spins, giving rise to *entangled plaquette states* [178, 179].

4.1.2 Path-integral Monte Carlo (PIMC)

> *At the heart of quantum mechanics is a rule that sometimes governs politicians or CEOs – as long as no one is watching, anything goes.*
>
> — Lawrence M. Krauss

To design a quantum-Monte-Carlo algorithm that does not rely on carefully designed wave-function ansätze, we turn our attention to simulating the quantum partition function $\mathcal{Z} = \operatorname{tr} e^{-\beta H}$, from which all physical observables (such as correlation functions) can be obtained at temperature $T = \beta^{-1}$. By breaking the exponential into many infinitesimal imaginary-time steps, $\mathcal{Z}$ can be represented as

$$\mathcal{Z} = \operatorname{tr}\left[\left(e^{-\varepsilon H}\right)^n\right] \approx \operatorname{tr}\left[(\mathbb{1} - \varepsilon H)^n\right] = \sum_{\boldsymbol{\sigma}_1,\ldots,\boldsymbol{\sigma}_n} \langle \boldsymbol{\sigma}_1|\mathbb{1} - \varepsilon H|\boldsymbol{\sigma}_2\rangle \cdots \langle \boldsymbol{\sigma}_n|\mathbb{1} - \varepsilon H|\boldsymbol{\sigma}_1\rangle, \quad (4.6)$$

where $\varepsilon \to 0$ and $n \to \infty$ such that $\beta = n\varepsilon$. Likewise, the expectation value of any operator A that is diagonal[†] in the computational basis, $A = \sum_{\boldsymbol{\sigma}} A(\boldsymbol{\sigma})|\boldsymbol{\sigma}\rangle\langle\boldsymbol{\sigma}|$, is

$$\langle A\rangle = \frac{1}{\mathcal{Z}} \operatorname{tr}\left(Ae^{-\beta H}\right) = \frac{1}{\mathcal{Z}} \sum_{\boldsymbol{\sigma}_1,\ldots,\boldsymbol{\sigma}_n} A(\boldsymbol{\sigma}_1)\langle \boldsymbol{\sigma}_1|\mathbb{1} - \varepsilon H|\boldsymbol{\sigma}_2\rangle \cdots \langle \boldsymbol{\sigma}_n|\mathbb{1} - \varepsilon H|\boldsymbol{\sigma}_1\rangle. \quad (4.7)$$

Now, if we sample ensembles of $\boldsymbol{\sigma}_1, \ldots, \boldsymbol{\sigma}_n$ using the product of matrix elements in (4.6) (normalised by their sum $\mathcal{Z}$) as a probability distribution, the expectation value $\langle A\rangle$ is given by the Monte-Carlo average of $A(\boldsymbol{\sigma}_1)$. Since $\boldsymbol{\sigma}_1, \ldots, \boldsymbol{\sigma}_n$ can be regarded as samples of the imaginary-time path-integral representation of $\mathcal{Z}$ [which is in fact a continuum limit of (4.6)], this method is known as *path-integral Monte Carlo*. It is an *unbiased* method in that it does not depend on ansätze external to the problem and is able to recover the exact behaviour of the system, limited only by the statistical noise of Monte-Carlo sampling.

In the simplest implementation of the algorithm, the imaginary-time span β is divided into a large but finite number of steps, and the finite list of $\boldsymbol{\sigma}_1, \ldots, \boldsymbol{\sigma}_n$ is sampled

[†]Of course, it is possible to write down similar expressions for non-diagonal operators as well; we focus on the diagonal case only for simplicity.

using a standard importance-sampling technique. However, to minimise the error introduced by changing $e^{-\varepsilon H}$ to $\mathbb{1} - \varepsilon H$, ε has to remain small, which leads to an impractically large number n of steps for low-temperature simulations. This has substantial memory cost; furthermore, most samples $\boldsymbol{\sigma}_i$ will be identical to their neighbours and introducing an off-diagonal matrix element[†] of $\mathbb{1} - \varepsilon H$ becomes so unlikely that simple Metropolis–Hastings sampling eventually fails.

These problems can be solved by only keeping track of the imaginary times at which off-diagonal matrix elements are introduced [180, 181]. Between these events, only the diagonal matrix elements of the Hamiltonian act, so the time evolution (4.6) can easily be re-exponentiated, eliminating the discretisation error. Furthermore, new off-diagonal matrix elements[‡] can be introduced at any imaginary time according to a continuous probability distribution with a finite integrated probability, and so importance sampling remains viable.

As explained above, PIMC samples can be used to evaluate correlation functions without reference to explicit wave functions; by doing so at different imaginary times, the energy of the ground state (and sometimes even of excited states [98, 137, 182]) can be extracted. Furthermore, as the method is inherently finite-temperature (although ground-state properties can be approximated by taking $\beta \to \infty$), it gives access to thermal properties like the heat capacity: this is particularly difficult for variational methods that focus on the ground state only.

4.1.3 The sign problem in PIMC

> *Distinguishing the signal from the noise requires both scientific knowledge and self-knowledge: the serenity to accept the things we cannot predict, the courage to predict the things we can, and the wisdom to know the difference.*
>
> — Nate Silver, *The Signal and the Noise*

The key shortcoming of path-integral Monte Carlo (and the reason one can still write a chapter in a doctoral thesis about new numerical methods for quantum mechanics) is that the probability distribution in (4.7) is only well-defined for a special class of Hamiltonians. In particular, consider a Hamiltonian where some off-diagonal matrix elements[§]

[†] These are $O(\varepsilon)$ compared to the diagonal matrix elements that are close to 1.

[‡] Or rather, pairs of them, since the cyclicity of the trace (4.6) requires that any changes to the spin configuration $\boldsymbol{\sigma}_1$ be undone before reaching $\boldsymbol{\sigma}_1$ at the other end.

[§] There is no such limitation for diagonal matrix elements (except that they be real as H is Hermitian), because the corresponding matrix element of $\mathbb{1} - \varepsilon H$ will be positive in any case.

in the computational basis are positive or have an imaginary part: the corresponding matrix elements in $\mathbb{1} - \varepsilon H$ are not positive, and so the products of matrix elements in (4.6) cannot be guaranteed to be all positive, either. If such is the case, the "probability distribution" PIMC tries to sample is not a well-defined probability distribution, so direct sampling fails. This is called the *Monte-Carlo sign problem*.

A straightforward way around the sign problem is sampling a well-defined probability distribution proportional to the moduli of the matrix-element products, adding their phases by hand. This turns (4.7) into

$$\begin{aligned}\langle A\rangle &= \frac{\sum_{\boldsymbol{\sigma}_1,\dots,\boldsymbol{\sigma}_n} A(\boldsymbol{\sigma}_1)\langle\boldsymbol{\sigma}_1|\mathbb{1}-\varepsilon H|\boldsymbol{\sigma}_2\rangle\cdots\langle\boldsymbol{\sigma}_n|\mathbb{1}-\varepsilon H|\boldsymbol{\sigma}_1\rangle}{\sum_{\boldsymbol{\sigma}_1,\dots,\boldsymbol{\sigma}_n}\langle\boldsymbol{\sigma}_1|\mathbb{1}-\varepsilon H|\boldsymbol{\sigma}_2\rangle\cdots\langle\boldsymbol{\sigma}_n|\mathbb{1}-\varepsilon H|\boldsymbol{\sigma}_1\rangle} \\ &= \frac{\sum_{\boldsymbol{\sigma}_1,\dots,\boldsymbol{\sigma}_n} A(\boldsymbol{\sigma}_1)S(\{\boldsymbol{\sigma}\})\left|\langle\boldsymbol{\sigma}_1|\mathbb{1}-\varepsilon H|\boldsymbol{\sigma}_2\rangle\cdots\langle\boldsymbol{\sigma}_n|\mathbb{1}-\varepsilon H|\boldsymbol{\sigma}_1\rangle\right|}{\sum_{\boldsymbol{\sigma}_1,\dots,\boldsymbol{\sigma}_n} S(\{\boldsymbol{\sigma}\})\left|\langle\boldsymbol{\sigma}_1|\mathbb{1}-\varepsilon H|\boldsymbol{\sigma}_2\rangle\cdots\langle\boldsymbol{\sigma}_n|\mathbb{1}-\varepsilon H|\boldsymbol{\sigma}_1\rangle\right|} = \frac{\langle AS\rangle'}{\langle S\rangle'},\end{aligned} \tag{4.8}$$

where $S(\{\boldsymbol{\sigma}\})$ is the sign of $\langle\boldsymbol{\sigma}_1|\mathbb{1}-\varepsilon H|\boldsymbol{\sigma}_2\rangle\cdots\langle\boldsymbol{\sigma}_n|\mathbb{1}-\varepsilon H|\boldsymbol{\sigma}_1\rangle$ and the primed expectation values are with respect to the probability distribution introduced above. $\langle S\rangle'$ is called the *average sign*, and its modulus indicates the "badness" of the sign problem. If $|\langle S\rangle'| = 1$, all terms in the path integral have the same sign, so there is no sign problem; if it is much smaller than one, terms with positive and negative signs have an approximately equal weight in the probability distribution, leading to large cancellations in both the numerator and denominator of (4.8). Producing an accurate Monte-Carlo average in the latter case requires a huge number samples, and below a critical average sign, the signal will be lost in floating-point arithmetic errors.

Unfortunately, for any Hamiltonian with a sign problem, the average sign decays exponentially in both the system size and the inverse temperature β [183]: PIMC can only be used on these systems for very small system sizes at relatively high temperatures. Furthermore, analogous sign problems appear in a variety of other quantum-Monte-Carlo methods, such as Green's function Monte Carlo and stochastic series expansion [38, 184]. Indeed, finding the ground state of an arbitrary sign-problematic Hamiltonian *by any method* is NP-complete, that is, solving it efficiently (in polynomial time) would solve a host of other computationally hard problems, such as prime factorisation [185].

Many important problems in physics are free of the sign problem, including systems of bosons and of ferromagnetically interacting spins. Fermion systems, on the contrary, usually have a sign problem. Due to the exchange antisymmetry of fermion wave functions, the sign of matrix elements in such Hamiltonians depends on the labelling of basis states: apart from special cases, there is no consistent labelling that makes all of them negative. For this problem, many approaches have been proposed to alleviate and cure the sign problem, see, for instance, Refs. [38, 186–189].

The Marshall transformation

In the rest of this chapter, I am going to focus on the sign problem in antiferromagnetic Heisenberg models

$$H = \sum_{ij} J_{ij} \vec{\sigma}_i \cdot \vec{\sigma}_j = \sum_{ij} J_{ij} \left[\sigma_i^z \sigma_j^z + \tfrac{1}{2}(\sigma_i^+ \sigma_j^- + \sigma_i^- \sigma_j^+) \right], \tag{4.9}$$

where some of the J_{ij} are positive: the $J\sigma_i^+\sigma_j^-$ terms produce positive off-diagonal matrix elements in the σ^z computational basis whenever $J > 0$.

In general, these matrix elements give rise to a serious sign problem that renders PIMC inapplicable. An exception is the case of antiferromagnets on bipartite lattices: consider a lattice where the sites can be divided into two disjoint subsets, A and B, such that $J_{ij} \geq 0$ if i and j belong to different subsets and $J_{ij} \leq 0$ if they belong to the same one. Now, let us rotate all spins in subset A by π around the σ^z axis:

$$\sigma^z \to \sigma^z; \qquad \sigma^x \to -\sigma^x; \qquad \sigma^y \to -\sigma^y; \qquad \sigma^\pm \to -\sigma^\pm. \tag{4.10}$$

Since this is done to precisely one of the spins involved in each antiferromagnetic term of (4.9), all positive off-diagonal matrix elements $J\sigma_i^+\sigma_j^-$ turn into $-J\sigma_i^+\sigma_j^-$, removing the sign problem.[†] This is known as the *Marshall transformation* [177]; it reflects the lack of competing interactions in bipartite lattice antiferromagnets and their resulting ability to form simple ordered phases. By contrast, frustrated antiferromagnets live on non-bipartite lattices, and thus do not admit similar sign rules that would remove their sign problem.

4.1.4 The sign problem in variational approaches

The sign problem, as introduced in the previous section, is a property of the Hamiltonian, without any obvious consequences for ground-state wave functions or variational approximations thereof. However, the sign structures of a Hamiltonian matrix and its ground state are related by a corollary of the Perron–Frobenius theorem [190, 191]:

> *Every real symmetric matrix with non-positive off-diagonal elements has an eigenvector corresponding to its lowest eigenvalue with only nonnegative entries.*

An intuitive way of thinking about this result is that imaginary-time evolution under a sign-problem-free Hamiltonian causes only constructive interference, which results

[†] In fact, the original Hamiltonian had no sign problem in the sense of (4.8), either. Any string of off-diagonal matrix elements in the path integral has to take the configuration $\boldsymbol{\sigma}_1$ back into itself: on a bipartite lattice, this is only possible with an even number of antiferromagnetic matrix elements.

in a wave function with all positive entries in the limit of large imaginary time. For a sign-problematic Hamiltonian, by contrast, the same evolution gives rise to destructive interferences and a ground state with a nontrivial sign structure. For some variational ansätze, these interferences might also lead to large cancellations in evaluating the wave function $\Psi(\boldsymbol{\sigma};\theta)$, to a similar detrimental effect as a small average sign in (4.8). For instance, this is a nontrivial concern for tensor network approaches to sign-problematic Hamiltonians, requiring sophisticated strategies for contracting the tensor networks to avoid cancellations [192]. The sign problem also affects variational optimisation protocols: stochastic reconfiguration was found less vulnerable to it than plain gradient descent [168].

The success of mean-field wave functions for fermionic and antiferromagnetic systems stems in part from how they automatically impose a physically reasonable sign structure for all values of their variational parameters (namely, the terms of the mean-field Hamiltonian). This allows them to converge reliably to low-energy variational states, and to provide a good basis for Jastrow factors (4.4) that usually only change the wave-function amplitudes. Mean-field wave functions are also common starting points for *fixed-node diffusion Monte Carlo (DMC),* in which wave-function amplitudes are represented in an unbiased way, but the sign structure is provided as an ansatz to work around the sign problem [189, 193–195].

For antiferromagnets on bipartite lattices, the Marshall transformation (4.10) eliminates the sign problem of the Hamiltonian and thus makes all entries of the ground-state eigenvector positive. It follows that $\langle\boldsymbol{\sigma}|\prod_{i\in A}\sigma_i^z|\mathrm{GS}\rangle$ is positive for all computational basis states $\boldsymbol{\sigma}$ (the σ_i^z are the rotation operators that implement the Marshall transformation). Therefore, each wave-function entry $\langle\boldsymbol{\sigma}|\mathrm{GS}\rangle$ is positive (negative) if there is an even (odd) number of down spins in sublattice A: this is known as the *Marshall sign rule (MSR)* [177]. Let me note that this sign rule is easily implemented in the Huse–Elser ansatz (4.5) by setting

$$\operatorname{Im} h_i = \begin{cases} \pi/2 & i \in \text{sublattice } A \\ 0 & i \in \text{sublattice } B; \end{cases} \qquad \operatorname{Im} J_{ij} = 0. \tag{4.11}$$

Finally, the converse of the Perron–Frobenius theorem is not true: Hamiltonians with a mild sign problem can well have ground states with all positive entries [196, 197]. In practice, this is related to the presence of ordered phases in the phase diagram of frustrated antiferromagnets: as long as a Hamiltonian remains in the same phase as a trivial unfrustrated point, its ground-state sign structure remains similar to the sign rule for the latter, allowing variational ansätze to capture it (see §4.2.3 and Ref. 198). This also

helps explain the success of the Huse–Elser ansatz for the triangular lattice Heisenberg antiferromagnet [176], which, while frustrated, shows a classical 120° order.

4.2 Neural quantum states (NQS)

> *Although I do not suppose that either of us knows anything really beautiful and good, I am better off than he is – for he knows nothing, and thinks that he knows. I neither know nor think that I know.*
>
> — Socrates, as quoted in Plato's *Apology*

Deep neural networks have been the workhorse of the machine learning and artificial intelligence revolution of the last decade [164, 199]. Even though they have been invented long ago [200], the recent increase in available computational power and training data brought out their advantage over other machine learning techniques, first on the task of image classification [201]. Since then, sophisticated neural-network-based machine-learning approaches solved seemingly insurmountable computational challenges, such as defeating the best human players in chess and go, or recognising hundreds of millions of faces. This success derives from the ability of deep neural networks to represent arbitrarily complex functions accurately [202] and to "learn" them from a sufficiently large amount of data.

Variational ansätze for VMC must be flexible enough to represent and readily converge to the complex pattern of probability amplitudes that arise in the ground states of challenging systems, such as quantum spin liquids. This task is similar to that of *unsupervised learning*, where neural networks are used to represent and sample probability distributions consistent with training data (e.g., given many pictures of humans, the neural network is to produce realistic pictures of humans that are different from those it was trained on). The success of neural networks on these tasks is a clear indication that they are able to represent highly nontrivial probability distributions on very high-dimensional spaces; therefore, it is natural to expect that they are able to represent and learn many-body wave functions equally well.

Neural networks (specifically, *restricted Boltzmann machines*) were first proposed as wave-function ansätze for VMC in 2017 [203]. Since then, a great multitude of *neural quantum state (NQS)* architectures have been proposed, utilising, among others, deep feed-forward neural networks [198, 204–209], recurrent neural networks [210–212], and Gaussian processes [213]. These have been deployed with promising results to challenging problems in electronic quantum chemistry [206–208, 214] and quantum spin liq-

uids [175, 204, 205, 215–217]; the ability of neural networks to represent arbitrary function without constraints like the entanglement limitations of tensor networks [166, 218] makes this approach desirable for higher-dimensional, critical, and topologically ordered systems [219, 220].

The key difference between representing probability distributions for unsupervised learning and wave functions for VMC is that the latter consists of probability *amplitudes,* with nontrivial signs or complex phases. These phases are precisely the origin of the Monte-Carlo sign problem, suggesting that neural quantum states may be susceptible to it. Indeed, representing highly nontrivial sign structures in, for instance, quantum-spin-liquid states has proven to be a serious challenge for NQS approaches [198, 205].

In the rest of this section, I shall review the first and still most popular NQS ansatz, restricted Boltzmann machines; NQS architectures based on deep neural networks; and the manifestations of the sign problem in NQS.

4.2.1 Restricted Boltzmann machines (RBM)

Restricted Boltzmann machines have first appeared in machine-learning literature as a physically motivated and flexible ansatz to model probability distributions [221]. They are defined as Ising models on a system of "visible spins" v_i (the degrees of freedom to be modelled) and "hidden spins" h_j: to make the resulting probability distribution tractable, we do not allow interactions between two visible or two hidden spins (hence "restricted"). The RBM probability distribution is the marginal probability of visible spins traced over the hidden spins:

$$p(\mathbf{v}) \propto \sum_{\mathbf{h}} \exp\left(a_i v_i + b_j h_j + w_{ij} v_i h_j\right) \propto e^{a_i v_i} \prod_j \cosh\left(b_j + w_{ij} v_i\right). \tag{4.12}$$

Here, summation over the indices i and j is implied, w_{ij} is the reduced interaction strength between the spins v_i and h_j, and a_i and b_j are "external fields" coupled to the same.[†] In the second equality, we explicitly sum over each $h_j = \pm 1$, made possible by the lack of interaction between hidden spins. It can be shown that an RBM with sufficiently many hidden spins can represent any probability distribution with arbitrary accuracy [222]. Furthermore, the distribution (4.12) can be sampled efficiently using *Gibbs sampling* [223]:

1. Given a visible spin configuration $\mathbf{v}$, sample a hidden spin configuration $\mathbf{h}$ from

[†]In machine-learning parlance, the w are called "weights," the a and b, "biases."

the conditional distribution

$$p(\mathbf{h}|\mathbf{v}) \propto \exp\left(b_j h_j + w_{ij} v_i h_j\right) = \prod_j \exp\left(b_j h_j + w_{ij} v_i h_j\right). \tag{4.13}$$

As there are no interactions between hidden spins, this distribution factorises, so each h_j is trivial to sample.

2. Using this hidden spin configuration, sample a new visible spin configuration $\mathbf{v}'$. Again, the conditional distribution $p(\mathbf{v}'|\mathbf{h})$ factorises and can easily be sampled.

This sampling procedure has better mixing properties than simple Metropolis–Hastings sampling of the overall probability distribution (4.12). Nevertheless, with the advent of recurrent neural networks [224], which do not rely on Markov-chain Monte Carlo at all and return completely uncorrelated samples, RBMs have lost traction for cutting-edge machine-learning applications. Owing to their simple and interpretable structure, however, they remained popular in theoretical studies and as ingredients of more complex architectures [225].

Restricted Boltzmann machines have first been proposed as a wave-function ansatz by Ref. 203, and have since remained the most popular NQS architecture [220]. In this case, the probability *amplitude* is represented by the RBM form (4.12); to allow for non-positive amplitudes, the weights $\mathbf{w}$ and biases $\mathbf{a}, \mathbf{b}$ are allowed to be complex. In fact, the summation over hidden units is merely symbolic: since the probability distribution $p(\boldsymbol{\sigma})$ is the modulus squared of the RBM output, there is no consistent notion of the joint probability distribution $p(\mathbf{v}, \mathbf{h})$, so Gibbs sampling becomes impossible. In practice, this is not a problem, as similar basis states tend to have amplitudes of the same order of magnitude, so Metropolis–Hastings sampling remains viable.

4.2.2 Deep (convolutional) NQS

The visual representation of an RBM suggests an analogy to single-layer (shallow) neural networks [Fig. 4.1(a)]. It is, therefore, a natural idea to enhance their expressive power by turning to *feed-forward deep neural networks.* These are compositions of alternating linear maps and nonlinear *activation functions:*

$$f(\mathbf{x}) = \sigma(\mathbf{W}_n \sigma(\mathbf{W}_{n-1} \sigma(\cdots \sigma(\mathbf{W}_1 \mathbf{x} + \mathbf{b}_1) \cdots) + \mathbf{b}_{n-1}) + \mathbf{b}_n), \tag{4.14}$$

where the weights $\mathbf{W}_i$ are matrices, the input $\mathbf{x}$ and biases $\mathbf{b}_i$ are vectors, and the nonlinear function $\sigma(\cdot)$ is understood to act elementwise. Such neural networks are often represented as a graph, where nodes, arranged in layers, stand for applications of the

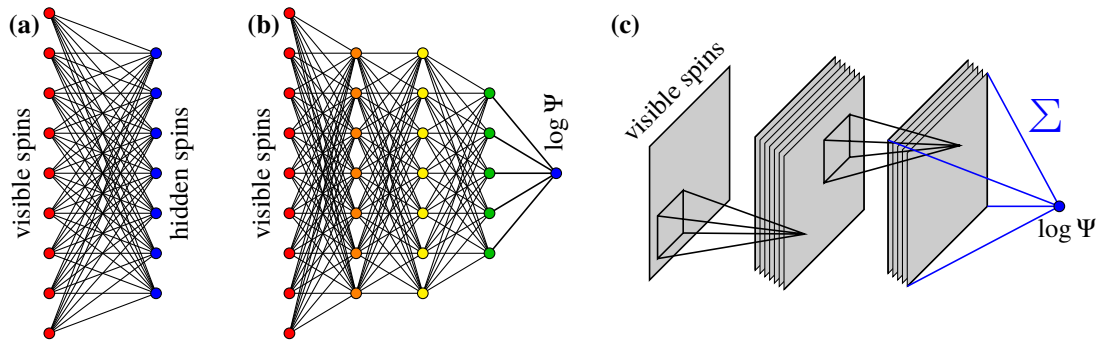

Figure 4.1. Common neural quantum state ansätze.
(a) A restricted Boltzmann machine consists of a number of "hidden spins" coupled to the physical ("visible") spins via Ising interactions: the wave function amplitude is given by the probability distribution of this Ising model, (4.12). It can be regarded as a single-layer feed-forward neural network.
(b) A fully connected deep feed-forward neural network consists of layers of hidden units, connected by unrestricted linear maps and nonlinear activation functions. To account for the multiplicatve nature of the wave function, the output of the last layer is taken to represent the logarithm of $\langle \boldsymbol{\sigma} | \Psi \rangle$.
(c) A convolutional neural network is a feed-forward network in which hidden units in each layer are arranged as replicas of the geometry of the input, and the linear map connecting the layers is a convolution with some kernel $w(\vec{r}' - \vec{r})$, resulting in a translation-invariant output. Often, the range of the kernels is restricted to enforce locality.

activation function and the biases $\mathbf{b}$, while edges represent the weights $\mathbf{W}$ [Fig. 4.1(b)]. Common activation functions include sigmoid functions [e.g., $\sigma(x) = \tanh x$] and rectified linear units [ReLU: $\sigma(x) = \max(x, 0)$].

An important feature of feed-forward neural networks is that derivatives of f with respect to the weights, biases, and inputs (all necessary for gradient descent optimisation) can be computed efficiently by iterating the chain rule. Consider a single layer of the network that maps its input $\mathbf{z}$ to the output $\mathbf{y} = \sigma(\mathbf{W}\mathbf{z} + \mathbf{b})$: given the derivatives $\partial f / \partial y_j$, we have

$$\frac{\partial f}{\partial z_i} = \sum_j W_{ij} \sigma'(y_j) \frac{\partial f}{\partial y_j}; \qquad \frac{\partial f}{\partial W_{ij}} = z_i \sigma'(y_j) \frac{\partial f}{\partial y_j}; \qquad \frac{\partial f}{\partial b_j} = \sigma'(y_j) \frac{\partial f}{\partial y_j}. \tag{4.15}$$

Using (4.15), the derivatives of f can be obtained working from the last layer backwards; therefore, this protocol is called *back-propagation* [226].

In general, a neural network can allow an arbitrary pattern of weights: these are called *fully connected networks*. However, these make no use of symmetries and other features of the problem at hand. A very important symmetry of both quantum lattice problems and many machine-learning tasks (e.g., image recognition) is translation in-

variance: a straightforward way of making use of this symmetry is arranging the layers beyond the input as replicas of the input geometry[†] and requiring that the weights and biases be translation invariant.[‡] Furthermore, weights can be constrained to be short-range (that is, each input would only affect nearby sites in the output): this enforces locality, another meaningful constraint in both machine-learning and quantum-mechanics contexts. Such architectures are called *convolutional neural networks* [Fig. 4.1(c)]: they have been extremely popular and successful in machine-learning tasks such as image recognition [164, 201]. The constraints of translation invariance and locality substantially reduce the number of variational parameters, making deeper and more expressive networks viable; furthermore, they help guide the optimisation of the network, stopping it from getting stuck at poor representations of the data that do not obey translation invariance [227].

Now, one can use any deep neural network architecture as a wave-function ansatz $\Psi(\boldsymbol{\sigma}; \mathbf{W}, \mathbf{b})$, using some convenient representation of the spins $\boldsymbol{\sigma}$ as the input $\mathbf{x}$ in (4.14). Since translation invariance is a crucial property of ordered lattice Hamiltonians and their ground states, I shall focus on convolutional networks here. Finally, it is important that neural networks of the form (4.14) are *additively extensive*, that is, the output of a juxtaposition two networks, not coupled by any of the weights, would be the sum of the outputs of the individual networks. This is at odds with many-body wave function amplitudes and multivariate probability distributions, that are usually *multiplicatively* extensive. The standard solution for this problem is having the neural network represent the logarithm of the wave function $\Psi(\boldsymbol{\sigma})$ [205, 228], which has the added benefit that back-propagation immediately produces the logarithmic derivatives O_j used in stochastic reconfiguration (4.3) [228]. In this formulation, RBM states are indeed single-layer neural networks with an activation function $\sigma(x) = \log\cosh x$, followed by summing all outputs.

[†] In machine-learning applications, *pooling layers* are also common: these reduce the size of the input image and thus force extracting relevant larger-scale features. For wave functions, maintaining exact translation invariance is desirable, so I shall not consider these.

[‡] That is, biases must be the same for all outputs in the same replica, and weights may only depend on the displacement between the input and the output site, not the sites themselves. In other words, the input is *convolved* with some kernel $w(\vec{r}' - \vec{r})$, hence the name "convolutional."

4.2.3 The sign problem in NQS

> *The thing that hath been, it is that which shall be; and that which is done is that which shall be done: and there is no new thing under the sun.*
>
> — Ecclesiastes 1:9

There is no *a priori* reason why NQS ansätze should suffer from the sign problem: neural networks are able to represent arbitrary functions [202, 222], including ones with positive and negative values, and stochastic reconfiguration is able to guide other ansätze (e.g., Jastrow–mean-field wave functions) to ground states with nontrivial signs [168]. In practice, however, neural-network wave functions have often struggled to attain state-of-the-art variational energies for sign-problematic Hamiltonians. This appears less pointedly in RBMs and other shallow, fully connected architectures, which are able to reach low variational energies even for Hamiltonians with severe sign problems [198, 204, 209, 217, 220, 229, 230], especially in cases where the sign problem can be cured by a Marshall-type sign rule [203]. By contrast, deeper networks and convolutional architectures of any depth quite often fail for antiferromagnetic or fermionic systems [205]: attempting to learn such states generates unphysically rough amplitude profiles, resulting in poor convergence or even complete breakdown of the protocol.

The most systematic computational approach to date to understand the origins of this manifestation of the sign problem has been offered by Ref. 198. In this work, the ground states of small (up to 36 spins) antiferromagnetic spin Hamiltonians on several different lattices were obtained by exact diagonalisation; various neural networks were then trained separately on the amplitudes and signs of a small subset of the entries in a supervised learning scenario (that is, the mismatch between network outputs and the exact wave function was minimised), and their ability to predict amplitudes or signs not used in the training was assessed. Since VMC algorithms rely on reconstructing operator expectation values from a very small sample of the Hilbert space, this *generalisation* ability is a useful proxy of how well the neural network performs as a wave-function ansatz in such applications. For every NQS architecture, sign structures proved harder to generalise than wave-function amplitudes. Furthermore, the quality of generalisation declined sharply upon entering spin-liquid phases: this reflects that the sign structures of spin liquids are much more complex than those of ordered magnets, so they are harder to reconstruct without additional physical insight.

As of now, there is no fully satisfactory solution of the NQS sign problem, especially for deep convolutional networks. For Heisenberg antiferromagnets on bipartite lattices, the Marshall sign rule can be used to explicitly cure the sign problem [203]. Surprisingly,

the same sign rule also stabilises VMC with deep networks for the square-lattice J_1–J_2 Heisenberg antiferromagnet in its spin-liquid phase, and allows it to attain state-of-the-art variational energies [205].

4.3 Alleviating the sign problem in NQS

> *Everyone shall be remembered, but each became great in proportion to his expectation. One became great by expecting the possible, another by expecting the eternal, but he who expected the impossible became greater than all.*
>
> — Søren Kierkegaard, *Fear and trembling*

The results of Ref. 198 suggest that the key limitation of NQS ansätze when it comes to approximating ground states of frustrated magnets is the ability of standard neural networks to represent the sign structure of the same states in a way that generalises reliably to basis states that were not used in the VMC protocol. Therefore, I focused on developing new NQS architectures specifically for sign structures in order to overcome the sign problem. I split the ansatz for $\log \Psi(\boldsymbol{\sigma})$ into two real-valued neural networks that encode the amplitude and phase of the wave function, respectively:

$$\log\langle \boldsymbol{\sigma}|\Psi\rangle = A(\boldsymbol{\sigma}) + i\Phi(\boldsymbol{\sigma}). \tag{4.16}$$

Standard deep, convolutional neural networks are able to represent the logarithmic amplitudes A accurately [205, 210]: this is to be expected, as the squares of these amplitudes give a probability distribution, which are captured excellently by deep neural networks (§4.2).

4.3.1 Representing the sign structure

As discussed in §4.2.2, the typical output of a neural network is the sum of a number of outputs in the last layer (in convolutional networks, there is at least as many of these as lattice sites). When it comes to Φ, each of these entries would correspond to a rotation on an Argand diagram of the wave function entry $\langle \boldsymbol{\sigma}|\Psi\rangle$ (Fig. 4.2, left): orchestrating these rotations to always yield an integer multiple of π is challenging in all cases but the trivial $\Phi \equiv 0$. This gives an intuitive picture of why out-of-the-box neural networks fail when it comes to representing nontrivial sign structures.

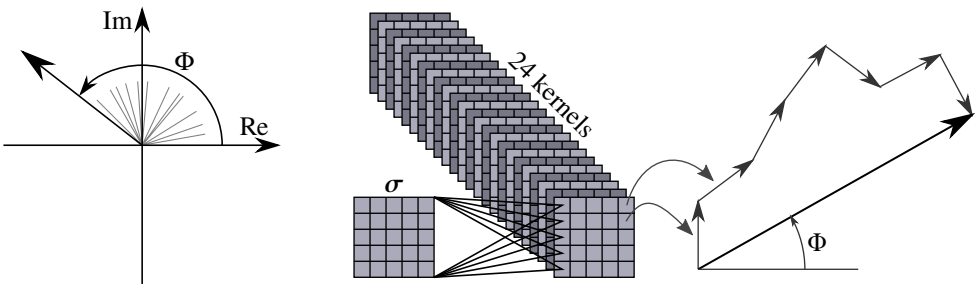

Figure 4.2. Left. In a typical deep network, the output is the sum of entries in the last layer: when representing phases, each of these entries (indicated by the angles between the grey lines) acts as an elemental rotation on an Argand diagram. **Mid-right.** Single-layer convolutional network used in our work to represent the phase Φ of the wave function. The spins $\boldsymbol{\sigma}$ are mapped through convolutional kernels spanning the entire lattice (with periodic boundary conditions). Each entry in the images is then taken as the argument of a unit complex number; Φ is the argument of their sum. Figure taken from Ref. 3.

Remaining at the mental image of the Argand diagram, a more natural way of combining phases is adding them as *phasors,* that is, complex numbers which carry the phases as their arguments (Fig. 4.2, right). This way, the overall phase Φ is more of an average than a sum of the individual phases, that is, if each underlying phase is able to converge approximately to the true phase, Φ will be a closer approximation of the same. It also respects the invariance of phases under adding multiples of 2π, a key mathematical property of phases that is completely ignored by typical neural networks.

The simplest full network that demonstrates this idea consists of a single convolutional layer, the outputs of which are summed as phasors (Fig. 4.2):

$$\Phi(\boldsymbol{\sigma}) = \arg\left[\sum_{n,\vec{r}} \exp\left(ib_n + \sum_{\vec{r}'} iw_{n,\vec{r}-\vec{r}'}\sigma^z_{\vec{r}'}\right)\right]. \tag{4.17}$$

where $\vec{r}, \vec{r}'$, and n label lattice sites and convolutional replicas, respectively; $w_{n,\vec{r}}$ and b_n are real-valued weights and uniform biases of the convolutional filters. The inner parenthesis convolves the spin configuration $\boldsymbol{\sigma}$ with the weights w; the imaginary exponentials of these convolutions are the phasors (of unit length) which are summed to obtain the final phase Φ. To allow for the greatest possible flexibility, the kernels w span the entire lattice in periodic boundary conditions.

4.3.2 Optimisation protocol

Even with adequate representational power, neural-network wave functions may struggle to converge to an antiferromagnetic ground state in VMC. Deep neural net-

works are typically initialised *critically*, that is, the typical magnitude of entries remains the same in each layer [231, 232]. It follows that the network outputs are typically small; since they represent the logarithms of the wave-function entries $\langle \boldsymbol{\sigma} | \Psi \rangle$, the latter will all be approximately 1. The state where all amplitudes are exactly equal is a classical product state in which all spins point along the $+\sigma^x$ direction; clearly, any state close to this one will have pronounced *ferromagnetic* correlations. Turning such a state into one with antiferromagnetic correlations is a highly unstable procedure, often leading to an unphysically rough amplitude profile: once this happens, Monte-Carlo sampling will either get stuck on a single $\boldsymbol{\sigma}$ or return extremely noisy data for stochastic reconfiguration, either of which makes further convergence impossible.

Separating the wave function ansatz into amplitudes and phases helps solve this problem too. The energy due to an antiferromagnetic matrix element $+J\sigma_i^+\sigma_j^-$ is minimised if the phases of the wave function entries coupled by it differ by π, regardless of their magnitudes: therefore, it is reasonable to expect that the optimal antiferromagnetic sign structure depends weakly on the amplitude profile.[†] One can, therefore, obtain a sign structure with broadly antiferromagnetic correlations by optimising $\Phi(\boldsymbol{\sigma})$ only with a fixed dummy amplitude profile $A(\boldsymbol{\sigma})$; the easiest is to set $A \equiv 0$, which makes all amplitudes equal. This sign structure is then used as a starting point for optimising both A and Φ together in standard stochastic reconfiguration, without risking severe instabilities in the former.

4.4 Numerical experiments

> *It doesn't matter how beautiful your theory is – if it doesn't agree with experiment, it's wrong.*
>
> — Richard Feynman

In this section, I demonstrate the effectiveness of the approach outlined in the previous section on the frustrated J_1–J_2 Heisenberg antiferromagnet (HAFM) on the square lattice:

$$H = J_1 \sum_{\langle ij \rangle} \vec{\sigma}_i \cdot \vec{\sigma}_j + J_2 \sum_{\langle\langle ij \rangle\rangle} \vec{\sigma}_i \cdot \vec{\sigma}_j, \tag{4.18}$$

where $J_1, J_2 \geq 0$ and $\langle ij \rangle$ and $\langle\langle ij \rangle\rangle$ refer to nearest and second-neighbour sites, respectively. This Hamiltonian is, in general, frustrated because of the triangles of antiferromagnetic interactions spanned by two nearest-neighbour and one second-neighbour bonds. However, it has two limits in which the frustration disappears (Fig 4.3):

[†] In unfrustrated antiferromagnets, all interactions can be optimised in this sense simultaneously, meaning that the Marshall sign rule is optimal for any set of variational amplitudes [177].

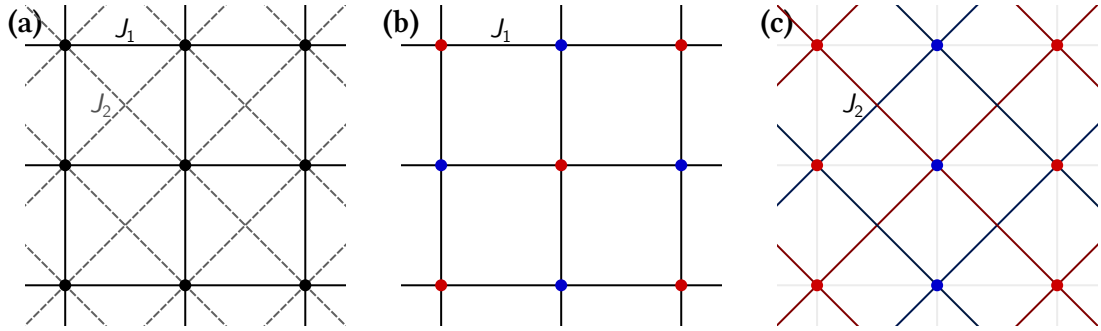

Figure 4.3. **(a)** The J_1–J_2 model on the square lattice is frustrated for general $J_1, J_2 > 0$.
(b) In the nearest-neighbour-only case, the square lattice is bipartite and splits into two chequerboard sublattices (red and blue dots), allowing us to define a Marshall sign rule using one of them. The ground state has a Néel order, where spin correlations follow a similar chequerboard pattern.
(c) If $J_1 = 0$, the square lattice decouples into the previous two chequerboard sublattices; J_2 acts as a nearest-neighbour coupling in the resulting two square lattices (maroon and dark blue edges). A Marshall sign rule can be defined on both of them, which readily combine into a stripy pattern along either columns (shown here) or rows.

- If $J_2 = 0$, there are only nearest-neighbour interactions: these are not frustrated as the square lattice is bipartite. The sublattices A and B in the Marshall transformation (4.10) follow a chequerboard pattern [Fig. 4.3(b)]. The corresponding ground state is *Néel ordered,* that is, positive and negative spin correlations also form a chequerboard pattern.

- If $J_1 = 0$, the lattice decouples into two chequerboard sublattices, each of which becomes a square lattice with diagonal axes and "nearest-neighbour" interactions [Fig. 4.3(c)]. A Marshall sign rule using alternating rows or columns can cure the sign problem of each half. Accordingly, the ground state shows *stripy order* with positive and negative spin correlations alternating in rows or columns.

Both unfrustrated points give rise to an ordered phase for finite values of J_2/J_1. Between these, one or more phases without long-range order emerge at $J_2/J_1 \approx 0.5$; determining their precise nature has been a long-standing challenge [174, 175, 196, 197, 233–237]. Recent high-quality evidence [174, 175] suggests a phase diagram consisting of Néel ordered, gapless spin-liquid, valence-bond solid, and stripy ordered phases (Fig. 4.4). This complex phase diagram, featuring both simple ordered and highly nontrivial fluctuating phases, has made the J_1–J_2 model a popular benchmark for new numerical techniques [205, 215–217]

We tested our protocol at the unfrustrated point $J_2 = 0$ and inside the gapless spin-liquid phase at $J_2/J_1 = 0.5$. We considered a 10×10 lattice with periodic boundary

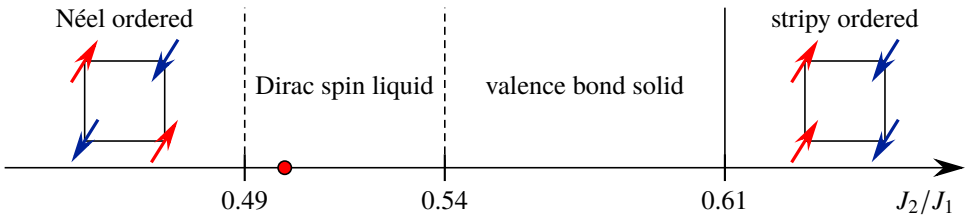

Figure 4.4. Phase diagram of the J_1–J_2 model (4.18) in the regime $J_1, J_2 > 0$, based on Ref. 175. Solid and dashed lines denote first and second order phase transitions, respectively. The red dot marks $J_2/J_1 = 0.5$.

conditions: this system size is substantially larger than what is tractable using exact diagonalisation, but we could compare our results to energy estimates obtained using a variety of state-of-the-art methods [205, 215, 216, 236, 238]. We set $J_1 = 1$ without loss of generality.

4.4.1 Details of implementation

The NQS ansätze were implemented and the VMC optimisation was performed using NetKet [228], a recently published machine learning library customised for quantum many-body problems. It provides an easy-to-use implementation of Hamiltonians and other operators, as well as stochastic reconfiguration. I extended the library with several ingredients for the ansatz proposed in §4.3.1, e.g., the phasor sum "activation layer," a split amplitude–phase network architecture, and several activation layers.

Neural network architectures

The amplitude structure $A(\boldsymbol{\sigma})$ was represented using a six-layer convolutional neural network in all numerical experiments. The layers consist of 8, 7, 6, 5, 4, and 3 10×10 lattice replicas, respectively. These are connected by convolutional filters with real-valued kernels spanning 4×4 sites in periodic boundary conditions. The ReLU activation function $\sigma(x) = \max(x, 0)$ was used in all layers except the last one. There, the wave-function amplitude was taken to be the modulus of the product of all entries in the last convolutional layer: this was achieved using a final activation function $\sigma_6(x) = \ln|x|$, followed by summing all entries. All convolutional layers before the last one are initialised with Gaussian distributed random numbers of zero mean and standard deviation chosen so as to preserve the typical magnitude of backpropagated derivatives [232]. The last set of kernels are initialised with a uniform bias of 1.0 and Gaussian distributed kernel entries with standard deviation 2×10^{-4}. This results in amplitudes uniformly close to 1 upon initialisation.

We employed 24 lattice replicas in the phasor sum ansatz (4.17) (cf. Fig. 4.2). Similarly good results are achieved using fewer replicas, but the wider network allows for faster and more reliable convergence. The convolutional filters are initialised with Gaussian distributed random numbers of standard deviation 0.043.

Optimisation protocol

We optimised each wave function ansatz via stochastic reconfiguration (SR) in two stages. First, the phases are optimised with a uniform amplitude distribution (i.e., setting $A \equiv 0$): 10 000 such SR steps with learning rate $\eta = 0.01$ were followed by 10 000 steps with $\eta = 0.05$. Next, both A and Φ were optimised starting from the phase distribution achieved in the first stage: for this, we used 5 000 steps with $\eta = 0.001$, 5 000 steps with $\eta = 0.01$, and, finally, 50 000 steps with $\eta = 0.05$. The learning rate was increased during the optimisation because the imaginary-time evolution emulated by SR results in infinitesimal temperature (and thus energy) reduction close to the ground state. In both stages, the Monte-Carlo averages in (4.3) were evaluated using 5 000 samples obtained via the Metropolis–Hastings algorithm as implemented by NetKet [228]. To enforce $\mathrm{U}(1)$[†] spin-rotation symmetry around the S^z axis, only basis states with $\sum \sigma^z = 0$ were sampled.

We found that the first stage quickly attains an approximately constant minimum of variational energy in both the unfrustrated and the spin liquid case. However, any residual optimisation speeds up the next stage significantly, which is desirable as a single, simple neural network can be evaluated an order of magnitude faster than the full ansatz. Likewise, while the second stage reaches a nearly converged variational energy in about 20 000 steps, it is reduced slightly throughout the rest of the procedure, too.

4.4.2 Learning the Marshall sign rule at $J_2 = 0$

In the unfrustrated limit $J_2 = 0$, the sign problem is cured by the Marshall transformation (4.10); it follows that the corresponding Marshall sign rule (MSR) attains the lowest possible variational energy for any given amplitude profile $A(\boldsymbol{\sigma})$, including the uniform one used in the first stage. Therefore, the phase network $\Phi(\boldsymbol{\sigma})$ is expected to approximate

$$\Phi_{\mathrm{MSR}}(\vec{\sigma}) \equiv \pi \sum_{i \in A} \frac{1 - \sigma_i^z}{2} \pmod{2\pi} \tag{4.19}$$

by the end of the first stage of the optimisation protocol of §4.3.2. Fig. 4.5 shows that this is indeed the case: Φ is equal to Φ_{MSR} up to an irrelevant constant and small fluctuations.

[†]Enforcing the full $\mathrm{SU}(2)$ spin rotation symmetry group is by and large still an open challenge [239].

	$J_2 = 0$	$J_2/J_1 = 0.5$
our work	−0.671275(5)	−0.494757(12)
GWF [236]	−0.66935(1)	−0.49439(1)
GWF + RBM [216]	−0.67111(2)	−0.49575(3)
CNN [205]	−0.67135(1)	−0.49516(1)
best [178, 236]	−0.671549(4)	−0.49755(1)

Table 4.1. Variational energies (in units of J_1 per spin) attained in this work compared to other state-of-the-art energies on the same system. Our approach consistently outperforms plain Gutzwiller-projected mean-field wave functions (GWF) [236] and achieves similar accuracy to the RBM-enhanced GWF of Ref. 216 and the convolutional networks of Ref. 205 (CNN). The "best" energy is obtained using numerically exact stochastic series expansion for $J_2 = 0$ [178] and Lanczos-corrected GWF for $J_2/J_1 = 0.5$ [236]. Table taken from Ref. 3.

The distribution of $\Phi - \Phi_{\text{MSR}}$ becomes even sharper in the second stage of the optimisation, indicating a continuing improvement of the phase ansatz as the amplitudes are optimised.

Variational energies

We then optimised the amplitudes and phases of the NQS ansatz simultaneously. The evolution of variational energy throughout the procedure is shown by the blue curve in Fig. 4.6. In the end, we achieved a variational energy of −0.671275 per spin, 2.7×10^{-4} higher than the numerically exact energy given by stochastic series expansion [178]. This energy is only slightly above the one attained by the convolutional network of Ref. 205, even though the latter has substantially more variational parameters (7676 compared to our 5145) and its sign problem is cured by imposing the exact MSR (§4.2.3). Our method also outperforms Gutzwiller-projected fermionic mean-field wave functions [236], as well as ansätze using RBMs as a Jastrow-type many-body correlator on top of the same [216]: once the Marshall sign rule is learned accurately by our architecture, the flexibility of deep networks in describing complex probability distributions is favourable over the highly constrained mean-field ansatz. Our variational energies, together with relevant benchmarks, are summarised in Table 4.1.

Symmetry eigenvalues and other observables

To further assess the quality of our variational wave function, we evaluated the expectation value of the total spin operator $\langle \vec{S}^2 \rangle$ and the parity operator $\mathcal{P} = \prod \sigma^x$, as well as the statistical weight of the five irreducible representations (irreps) of the point group D_4 of the square lattice in the wave function. Since $\vec{S}^2$, $\mathcal{P}$, and the point-group

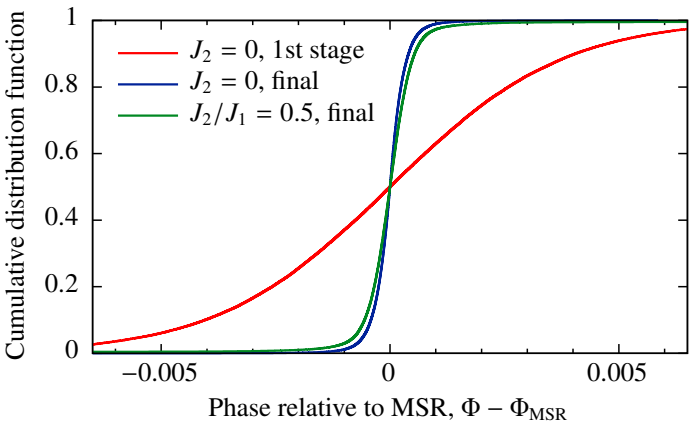

Figure 4.5. Cumulative distribution function of Marshall-adjusted phases $\Phi - \Phi_{\mathrm{MSR}}$ for the unfrustrated case $J_2 = 0$, learned before (red) and after (blue) optimising the amplitudes, and for the final state in the frustrated case $J_2/J_1 = 0.5$ (green). All distributions are very sharply peaked, indicating that Φ is a good approximation of the MSR (4.19). The standard deviations of the distributions are, in order, 3.3×10^{-3}, 4.0×10^{-4}, and 2.8×10^{-3}. Figure taken from Ref. 3.

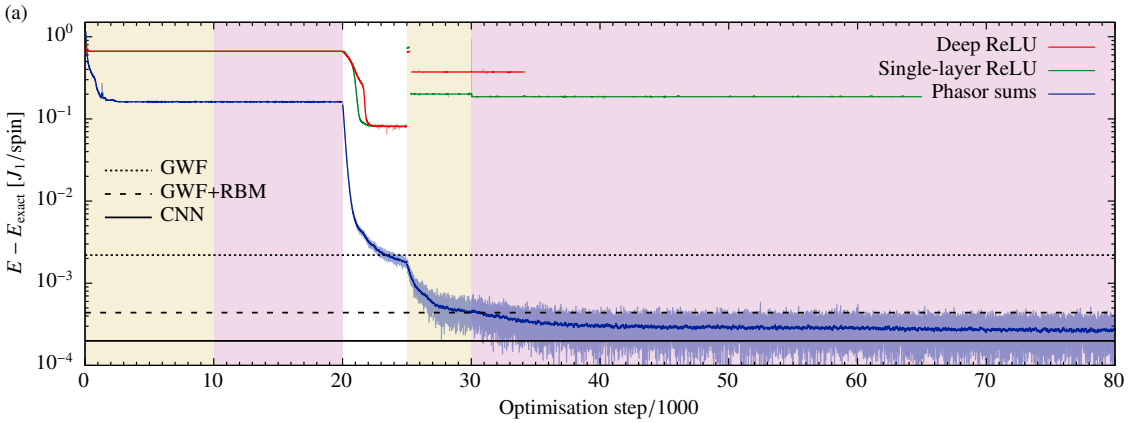

Figure 4.6. Convergence of the optimisation scheme of §4.4.1 with various neural-network ansätze (§4.4.1, §4.4.3) for the nearest-neighbour square-lattice HAFM. The shaded areas show the full spread of energy estimates used by the SR algorithm, the thicker lines show 100-step moving averages. The background shading indicates the learning rate η (white: 0.001, yellow: 0.01, purple: 0.05). The phasor sum ansatz (4.17) (blue curve) converges reliably to energies close to the true ground state; other ansätze (red and green curves; see §4.4.3 for details), both shallow and deep, fail to reach either a consistent variational energy, or one close to the ground state. Energies are compared to the numerically exact ground-state energy obtained by sign-problem-free stochastic series expansion [178]. For reference, variational energies are also shown for the convolutional network of Ref. 205 (CNN), plain [236], and RBM-improved [216] Gutzwiller-projected mean-field wave functions (GWF). Figure taken from Ref. 3.

	$J_2 = 0$	$J_2/J_1 = 0.5$
Parity $\langle \mathcal{P} \rangle$	0.998373(29)	0.990426(69)
Weight of irrep $\boldsymbol{A_1}$	**0.998645(18)**	**0.989363(51)**
A_2	0.000142(6)	0.000928(15)
B_1	0.000283(8)	0.003335(29)
B_2	0.000167(6)	0.001169(17)
E	0.000763(14)	0.005205(36)
Total spin $\langle \vec{S}^2 \rangle$	0.065(21)	0.581(43)
Stripy o.p. $S^2(\pi, 0)$	0.00498(5)	0.00521(7)
Néel o.p. $S^2(\pi, \pi)$	0.1571(2)	0.0633(2)

Table 4.2. Average parity $\langle \mathcal{P} \rangle$, total spin $\langle \vec{S}^2 \rangle$, and antiferromagnetic order parameters of, and statistical weight of irreps of the point group D_4 in, the fully converged NQS wave functions in the unfrustrated limit $J_2 = 0$ and for $J_2/J_1 = 0.5$. Both wave functions are predominantly parity even and transform according to the trivial representation A_1 (bold); the weights of states with odd parity and/or different spatial symmetry are around 0.001 and 0.01 in the two cases. The converged $\langle \vec{S}^2 \rangle$ is similarly larger in the frustrated case, consistent with its worse energy convergence. Table taken from Ref. 3.

symmetry operators all commute with the Hamiltonian (4.18), the true ground state is an eigenstate of the former two, and transforms according to precisely one irrep of the latter; furthermore, ground states of HAFMs are normally singlets ($\vec{S}^2 = 0$) and thus have even parity:[†] deviations from these expectations can be used as a quantitative test of the converged wave functions.

These results are shown in Table 4.2; computational details are given in Appendix C.2. We achieve similar figures of $\langle \vec{S}^2 \rangle$ to those of Ref. 205, which, together with the very low weight of odd-parity states, suggests a small admixture of states with high spin quantum numbers.[‡] We also evaluated the antiferromagnetic order parameter

$$S^2(\vec{q}) = \frac{1}{N(N+2)} \sum_{i,j} \langle \vec{\sigma}_i \cdot \vec{\sigma}_j \rangle e^{i\vec{q}\cdot(\vec{r}_i - \vec{r}_j)} \tag{4.20}$$

for $\vec{q} = (\pi, 0)$ and (π, π), which correspond to stripy and Néel orders, respectively: these results are consistent with those plotted in Ref. 205.

Interpreting the phase ansatz

The simplicity of the phase representation (4.17) allows us to interpret its behaviour directly by studying the kernel entries $w_{n,\vec{r}}$. These are plotted for all 24 kernels in Fig. 4.7.

[†] Simultaneous eigenstates of $\vec{S}^2$ and $S^z = 0$ are eigenstates of $\mathcal{P}$ as well; even and odd spin quantum numbers s correspond to $\mathcal{P} = +1$ and -1, respectively.

[‡] $\langle \vec{S}^2 \rangle \approx 0.06$ could be consistent with a 97:3 mixture of singlet and triplet states; however, this would yield an average parity $\langle \mathcal{P} \rangle = 0.94$. We cannot rule out a large $s = 2$ admixture on these grounds, but it is more natural to assume contributions with a range of higher spin quantum numbers.

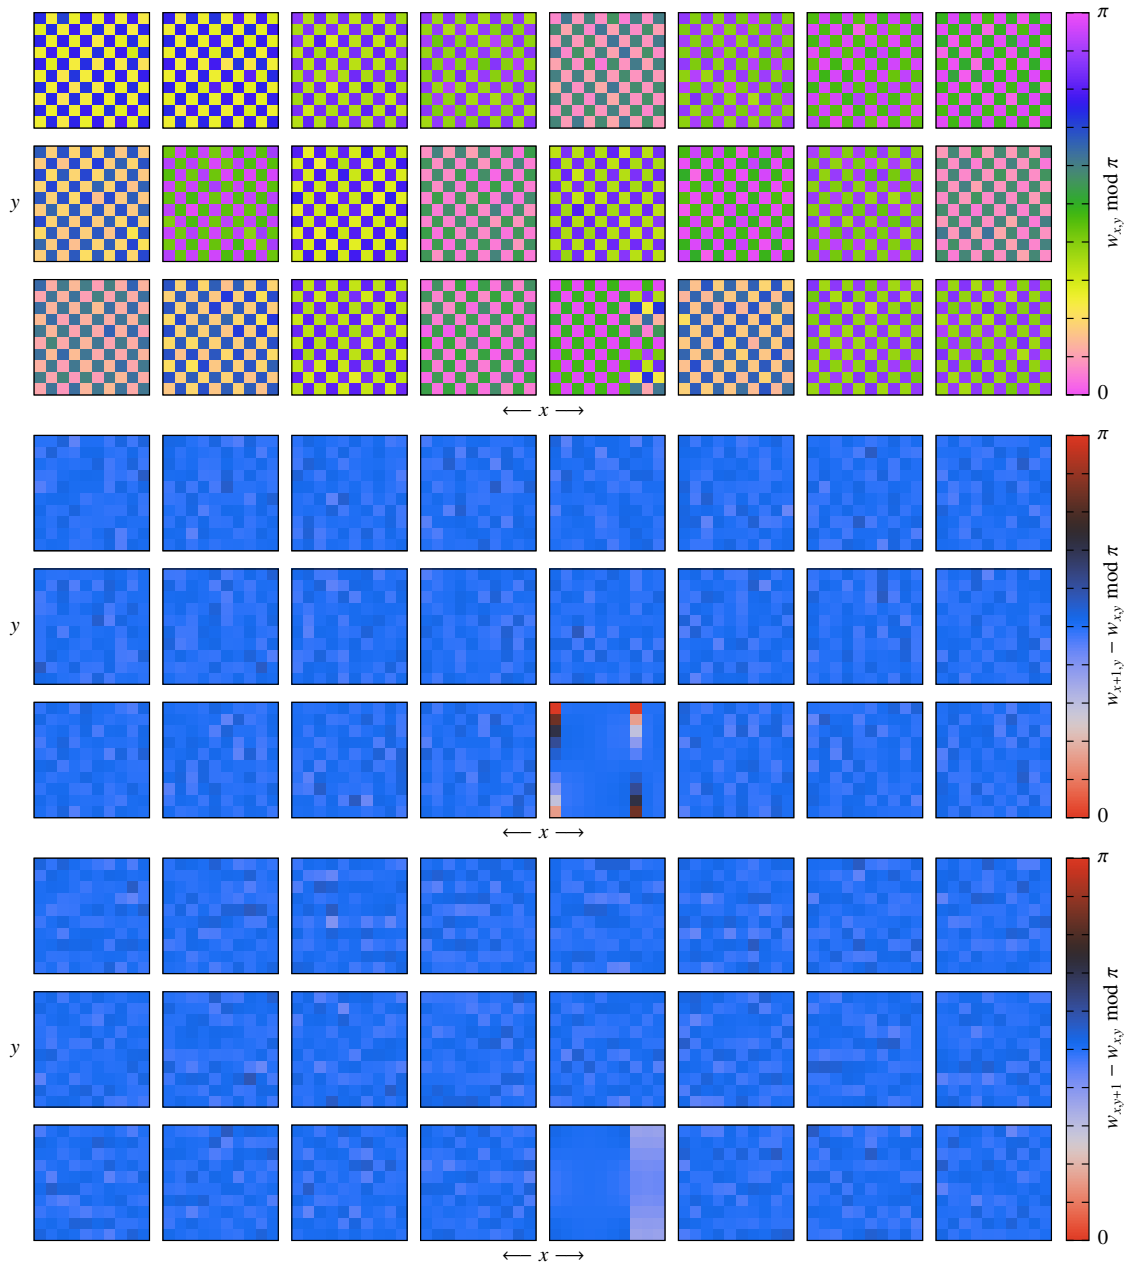

Figure 4.7. Weights $w_{n,\vec{r}}$ of all convolutional kernels n converged for $J_2 = 0$ (top three rows), and the differences $w_{n,\vec{r}+\hat{x}} - w_{n,\vec{r}}$ and $w_{n,\vec{r}+\hat{y}} - w_{n,\vec{r}}$ (middle and bottom three rows, respectively). All kernels but one show a clear chequerboard pattern, which results in $\Delta\phi \approx \pi$ upon exchanging a neighbouring up and down spin, consistent with the Marshall sign rule. In the only exception (fifth kernel of the third row), a "topological fault" spanning three columns appears, causing some $\Delta\phi$ to wind from 0 to 2π: these have little effect on the overall $\Delta\Phi$ and might persist due to a topologically invariant winding number. Figure taken from Ref. 3. Perceptionally uniform colour maps chosen following Ref. 125.

Almost all kernels display a distinct chequerboard pattern, with neighbouring entries differing by $\pi/2$ mod π; this is visually consistent with the chequerboard pattern of the Marshall sign rule in Fig. 4.3(b). To show that they are indeed equivalent, note that the MSR ensures that two basis states related by swapping nearest-neighbour up and down spins appear in the ground-state wave function with phases that differ by $\Delta\Phi = \pi$. In our ansatz (4.17), this $\Delta\Phi$ is an average of the changes $\Delta\phi$ of individual phasors in the sense that if all $\Delta\phi$ are the same, $\Delta\Phi$ will also be equal to them. Now, for swapping a pair of horizontal neighbour up and down spins, $\Delta\phi = 2(w_{\vec{r}+\hat{x}} - w_{\vec{r}})$, where $\hat{x}$ is the horizontal lattice vector and $\vec{r}$ ranges over the entire kernel:[†] in our chequerboard pattern, these are almost all close to π mod 2π (middle third of Fig. 4.7), so $\Delta\Phi \approx \pi$ to a good approximation. The same goes for vertical nearest neighbours as well (bottom third of Fig. 4.7).

In fact, fluctuations around the ideal $\Delta\phi = \pi$ tend to be self-averaging: as the kernel spans the entire lattice, $\Delta\phi$ must sum to 0 mod 2π around closed loops, so entries with $\Delta\phi \gtrsim \pi$ must be compensated by ones with $\Delta\phi \lesssim \pi$, which cancel out their effects on $\Delta\Phi$. That is, the perfect chequerboard pattern sits at the bottom of a wide, shallow basin in parameter space, making it easy to find for optimisation algorithms [164, 240].

The only surprising feature in Fig. 4.7 is a kernel that develops a three column wide "topological fault" in which the Marshall-adjusted convolutional weights wind around the vertical direction. This results in a number of $\Delta\phi$ far from the desired π. The large number of kernels, however, allows these to be corrected through slight deviations of the other kernels from the exact MSR. The additional kernels also play a key role in eliminating such detrimental structures: unwinding a "topological fault" requires large-scale changes in the individual phasors, which substantially increase the variational energy, unless corrected for by the other kernels. Indeed, our attempts to use a single convolutional kernel in (4.17) were plagued by robust "topological faults" spanning the entire kernel, leading to convergence far above the ground state energy.

4.4.3 Comparison of sign-structure ansätze

To demonstrate the advantage of our phase-structure ansatz over other convolutional networks, we considered two alternative architectures:

1. 24 convolutional filters spanning the entire lattice, followed by a ReLU activation layer and summation

2. The architecture used for the amplitudes, except for the last layer, where the $\ln|x|$ activation is also replaced with ReLU.

[†] The factor of two appears because up and down spins are represented by ± 1 rather than $\pm 1/2$.

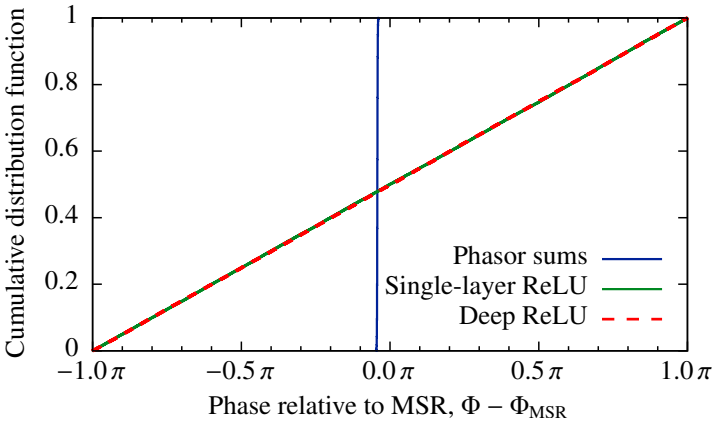

Figure 4.8. Cumulative distribution function of the Marshall-adjusted phases $\Phi - \Phi_{\text{MSR}}$ learned by the three phase structure ansätze in the first optimisation stage in the unfrustrated case $J_2 = 0$. The phasor-sum ansatz (4.17) develops a very narrow distribution, which reproduces the MSR to a good approximation; the other ansätze show an approximately uniform distribution, i.e., a complete failure to learn the sign rule. Figure taken from Ref. 3.

Amplitudes were encoded using the same ansatz as described in §4.4.1. The performance of these architectures under the protocol of §4.4.1 is shown by the green and red curves in Figs. 4.6 and 4.9, respectively. In the first stage, neither of them reaches the optimal variational energy found with the phasor-sum ansatz, indicating a failure to learn the MSR in the unfrustrated case. Subsequently, the amplitude network also fails to approach the ground state, even though it is capable of representing it closely (§4.4.2). We also point out that, as the Monte-Carlo sampling is restarted after changing the learning rate η, the estimates of the variational energy change substantially, which leads to discontinuities in Fig. 4.6. This indicates that the amplitude structure had developed several unphysically strong peaks, which make subsequent Monte-Carlo sampling unable to recover the correct wave function (cf. §4.2.3 and §4.3.2).

In the unfrustrated case, we also probe the phase structures learned by the various ansätze by comparing them to the exact Marshall sign rule directly. The distribution of the differences $\Phi - \Phi_{\text{MSR}}$ is shown in Fig. 4.8: while the phasor-sum architecture learns the MSR to a high accuracy (up to an irrelevant overall phase), the alternatives produce essentially random phases.

4.4.4 Inside the spin-liquid phase

At $J_2/J_1 = 0.5$, the model is inside a fully frustrated spin-liquid phase [175]. The wavefunction ansatz and optimisation protocol described in §4.4.1 still attains state-of-the-art variational energies, namely, -0.494757 per spin, 2.8×10^{-3} higher than the best energies

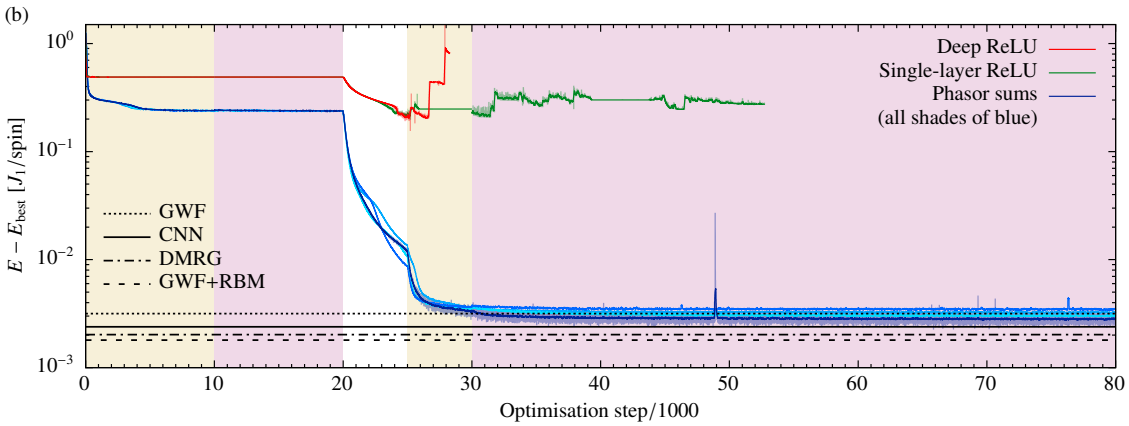

Figure 4.9. Convergence of the optimisation scheme of §4.4.1 with various neural network ansätze (§4.4.1, §4.4.3) for the square-lattice J_1–J_2 HAFM at $J_2/J_1 = 0.5$, inside the spin-liquid phase. The phasor-sum ansatz (4.17) (blue curves: 100-step moving averages for four runs; blue shading: full spread of energy estimates for the best run) outperforms the alternative phase architectures (red and green curves) again. The background shading indicates the learning rate η (white: 0.001, yellow: 0.01, purple: 0.05). Energies are compared to the best variational energy in the literature, obtained using Lanczos-corrected Gutzwiller-projected mean-field wave functions [236]. For reference, variational energies from Refs. [205, 216, 236] (cf. Fig. 4.6), as well as from DMRG [237] are also shown. Figure taken from Ref. 3.

obtained by Lanczos-iterating a Gutzwiller-projected mean-field wave function [236]. Again, our result compares favourably with the best NQS-based variational energy in the literature [216], where the corresponding error is 1.8×10^{-3}. Notably, the typical scale of energy discrepancies in all approaches is an order of magnitude higher than in the unfrustrated case. Our variational energy and relevant benchmarks are summarised again in Table 4.1.

Surprisingly, however, we find that the converged phase structure $\Phi(\boldsymbol{\sigma})$ in the frustrated case recovers the MSR to a high accuracy; no bimodality consistent with having both positive and negative amplitudes can be seen (Fig. 4.5). This is at odds with the fact that the frustrated Hamiltonian (4.18) remains sign-problematic even after the Marshall transformation (4.10), so its average sign should fall below one [185]. However, the MSR is expected to remain a relatively good approximation[†] of the exact ground-state sign structure (Appendix C.3), which explains the low variational energy attained with this incorrect sign structure, as well as the stabilising effect of the MSR on the out-of-the-box convolutional network of Ref. 205.

As in the unfrustrated case, we evaluated the total spin $\langle \vec{S}^2 \rangle$, the mean parity $\langle \mathcal{P} \rangle$,

[†]For smaller values of J_2/J_1, the MSR holds exactly [196, 197] (cf. §4.1.4), but the transition out of the Néel phase leads to a qualitative change in sign structure [198].

Néel and stripy antiferromagnetic order parameters, and the statistical weights of the five irreps of the D_4 point group (Table 4.2). We note that all discrepancies from the true ground state ($\vec{S}^2 = 0$; $\mathcal{P} = +1$; transforms according to a single irrep of D_4) are an order of magnitude higher compared to $J_2 = 0$: this is consistent with the worse energy convergence and sign structure discrepancies of the former. Our figure of $\langle \vec{S}^2 \rangle$ is again similar to that achieved in Ref. 205: this hints that the deep convolutional networks used in that work stuck to the MSR they were preconditioned with.

We studied the representation of the MSR attained in the spin-liquid case by plotting the entries $w_{n,\vec{r}}$ of all 24 kernels in Fig. 4.10. This representation is obviously more complex than the one in Fig. 4.7, with many "topological faults" and some kernels that show no discernible pattern. Many kernels, however, retain the chequerboard pattern consistent with MSR, and the $\Delta\phi$ upon exchanging nearest-neighbour spins vertically are clearly dominated by values close to π. Both of these are consistent with the fact that the kernels represent the MSR rather than the true frustrated sign structure; furthermore, as the latter is also close to the MSR (Appendix C.3), the emergence of such kernels would not be surprising in a more accurate representation of the ground state, either. We also observe several columns of the stripy pattern consistent with the MSR of the $J_1 = 0$ limit [Fig. 4.3(c)]. These result in $\Delta\phi \approx 0$ for horizontal nearest-neighbour exchanges, leading to a much more varied pattern in these differences. It is surprising that this more diverse structure has no apparent effect on the overall sign structure learned by the network.

Furthermore, the striking difference between $\Delta\phi$ along the horizontal and the vertical directions is not warranted either by any fundamental property of the ansatz, or by the converged wave function that obeys all point-group symmetries to a high accuracy (Table 4.2). Nevertheless, it might hint at spontaneous point-group symmetry breaking at higher variational energies that ultimately leads to learning an incorrect sign structure.

4.5 Outlook

> *It was rather too late in the day to set about being simple-minded and ignorant.*
>
> — Jane Austen, *Emma*

In summary, we developed a robust and efficient protocol for finding low-energy variational states with a nontrivial sign structure using convolutional neural quantum states without any prior knowledge on the sign problem of the Hamiltonian. We used an ansatz with two neural networks that represent amplitudes and phases separately

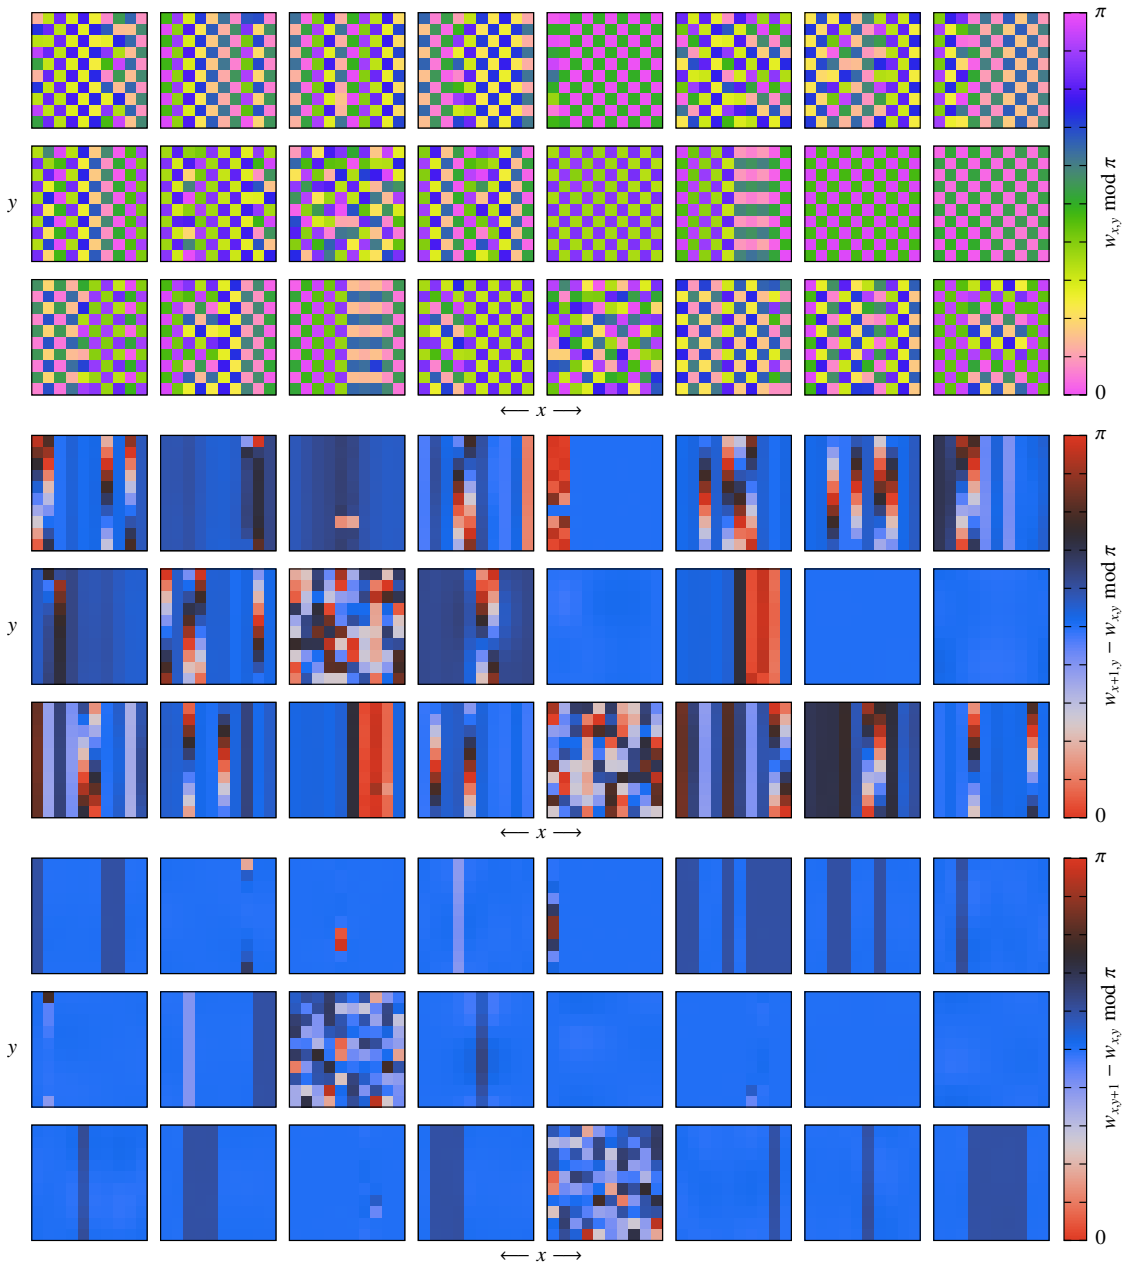

Figure 4.10. Weights $w_{n,\vec{r}}$ of all convolutional kernels n converged for $J_2/J_1 = 0.5$ (top three rows) and the differences $w_{n,\vec{r}+\hat{x}} - w_{n,\vec{r}}$ and $w_{n,\vec{r}+\hat{y}} - w_{n,\vec{r}}$ (middle and bottom three rows, respectively). The kernels are altogether much less regular than in the unfrustrated case, several of them completely scrambled and many showing various winding patterns. Nonetheless, the chequerboard pattern corresponding to the MSR is still common. In addition, a stripy pattern with successive rows differing by $\pi/2$ appears, leading to $\Delta\phi \approx 0$ upon exchanging nearest-neighbour up and down spins horizontally but $\Delta\phi \approx \pi$ vertically, apparently breaking rotational symmetry. Figure taken from Ref. 3. Perceptionally uniform colour maps chosen following Ref. 125.

and optimised it in two stages, first generating an approximate phase structure from which the entire wave function can readily converge without encountering severe instabilities. We demonstrated our approach by attempting to learn the ground states of the square-lattice spin-1/2 J_1–J_2 HAFM both at the unfrustrated point $J_2 = 0$ and at $J_2/J_1 = 0.5$, inside the fully frustrated spin-liquid phase. In both cases, we reached variational energies comparable to the best NQS energies reported in the literature [205, 216]; the difference might be attributed to the smaller number of variational parameters in our ansatz. Importantly, we used a fully convolutional architecture: this automatically imposes translation invariance, a useful inductive bias that speeds up convergence to a robust representation [227] and allows for resolving the lowest-energy states in different symmetry sectors [175, 209, 217]. Furthermore, the convolutional structure reduces the number of variational parameters from the $\mathcal{O}(N^2)$ typical for RBMs and other fully connected architectures [217, 241] to $\mathcal{O}(N)$, which keeps VMC algorithms viable for larger system sizes.

At $J_2 = 0$, our phase-structure ansatz (4.17) learns the Marshall sign rule with better generalisation properties than other convolutional networks, both deep and shallow, which is crucial for finding ground states reliably [198]. In the spin-liquid phase, however, the same approach fails to find the true ground-state sign structure, homing in instead on the same MSR. The existence of low-lying variational states with "simple" sign structures deep within frustrated phases parallels the poor generalisation of the corresponding ground states in supervised-learning scenarios [198], hinting at a possible bias of NQS ansätze towards "simple" states and the corresponding ubiquity of this "residual sign problem". While variational wave functions with sign structures that approximate the true ground state can improve the performance of algorithms like reptation Monte Carlo [242], they fall short of the ultimate goal of using NQS as accurate ground-state representations directly.

Going further from here raises important philosophical questions. Neural quantum states promise, in principle, a way of solving challenging quantum many-body problems with minimal input of physical insight, much like path-integral Monte Carlo. However, once out-of-the-box machine learning approaches do not deliver on this promise, we face some murky questions: *What is the least amount of physical insight that allows NQS to work? What* ***defines*** *what the least amount is? What* ***counts*** *as physical insight?* Different answers and attitudes to these questions can lead to drastically different strategies to overcome the sign problem and other limitations of NQS.

In this section, I discuss possible origins of the sign problem in out-of-the-box NQS approaches in terms of their functional form in the computational basis, without making reference to, e.g., fermionic mean-field wave functions. This allows me to suggest

improvements exploiting "mathematical insights" on the structure of typical wave functions: these improvements are largely based on mathematical features of tensor-network states and the Huse–Elser ansatz. Besides remaining closer to the original premise of NQS, such ansätze are more likely to perform well on generic systems, where physical insight on par with HAFMS (§1.2, §4.1.1) is unavailable. I also review two alternative strategies emerging in the latest literature: (i) employing sophisticated state-of-the-art neural network architectures; (ii) using neural networks to make mean-field ansätze more flexible.

4.5.1 Understanding the origin of the NQS sign problem

> *Never trust an experimental result until it has been confirmed by theory.*
>
> — Arthur Stanley Eddington

Thus far, I have referred to two different issues as the NQS sign problem: (i) the inability of any NQS ansatz to capture frustrated sign structures; (ii) the difficulties of (especially deep) neural networks to converge to ground states even with simple sign rules. I now discuss these "strong" and "weak" sign problems separately.

The strong sign problem: representing sign structures in frustrated phases

Ground states are the $\tau \to \infty$ limit of the imaginary-time evolution $e^{-\tau H}|\Psi_0\rangle$. This evolution, just like in real time, can be visualised as interferences between the computational basis states that make up the wave function. For sign-problem-free Hamiltonians, these interferences are always constructive (§4.1.4), leading to a wave function with only positive entries. Likewise, if the sign problem can be cured by a Marshall-type sign rule, no destructive interference arises as long as $|\Psi_0\rangle$ obeys the same rule. In these simple cases, homing in on the appropriate sign rule allows for focusing on optimising the amplitude profile, which is in turn not hampered by large cancellations.

The situation is more difficult for generic sign-problematic Hamiltonians (e.g., for fermionic or frustrated magnetic systems). There, no single sign structure is able to remove all destructive interferences from the imaginary-time evolution, which, even starting from a $|\Psi_0\rangle$ with the sign structure of the true ground state, may undergo temporary changes in signs. Mathematically speaking, tuning across $\Psi(\boldsymbol{\sigma}) = 0$ plays no special role in imaginary-time evolution under a generic Hamiltonian.

By contrast, typical NQS ansätze represent the wave function as the exponential of the (complex) neural network output (§4.2). As a result, the sign (or complex phase) of

the wave function is given by the imaginary part of the neural-network output: positive and negative amplitudes correspond to an imaginary part of 0 or π mod 2π, respectively. Changing the wave function $\Psi(\boldsymbol{\sigma})$ from a small positive to a small negative value in the process of VMC optimisation now requires changing the imaginary part of $\log \Psi(\boldsymbol{\sigma})$ by π: this is unlikely to be achieved by a continuous tuning of parameters, hampering the ability of the NQS ansatz to follow the exact imaginary-time evolution towards the ground state.

In general, it seems unlikely that ansätze that disregard the close connection between the evolution of amplitudes and sign structures in frustrated systems would succeed in representing their ground states. This explains the key observations of Ref. 198:

- in ordered phases, the sign structures are sufficiently close to a Marshall-type sign rule that even a shallow neural network can learn them reliably from relatively little training data;
- in spin-liquid phases, the sign structure is a fantastically complicated object without the context of the corresponding amplitudes; accordingly, a lot of training data is necessary to reconstruct it.

In recent literature, translation invariant ansätze have been implemented by adding wave-function estimates (rather than their logarithms) for translated copies of the system [175, 217]. These appear to outperform several of the NQS approaches described so far and can also represent and learn excited states in different wave-number sectors as

$$\Psi_{\vec{k}}(\boldsymbol{\sigma}) = \sum_{\vec{r}} \Psi'(\mathcal{T}_{\vec{r}}\boldsymbol{\sigma}) e^{i\vec{k}\cdot\vec{r}}, \tag{4.21}$$

where Ψ' is a non-translation-invariant wave function and $\mathcal{T}_{\vec{r}}$ stands for translation by the lattice vector $\vec{r}$.[†] Such an ansatz makes crossing $\Psi(\boldsymbol{\sigma}) = 0$ simple, as it only relies on cancellations between a small number of Ψ' outputs. The success of these ansätze is thus promising for the above interpretation of the strong sign problem.

The weak sign problem: representing simple sign rules with deep networks

This problem is harder to pinpoint, for it is much less universal. Several shallow and a few deep ansätze have achieved good variational energies for HAFMS on the bipartite square and honeycomb lattices, without explicitly using the Marshall sign rule. Furthermore, in the supervised-learning protocol of Ref. 198, sign structures in ordered phases

[†]Very recently, such constructions have been generalised to space-group symmetries and other non-Abelian groups using group convolutional networks [243, 244].

were learnt accurately by a range of shallow and deep, fully connected and convolutional architectures.[†]

The cleanest examples of the "weak sign problem" are seen in VMC using fully convolutional deep neural networks [3, 205]. I believe these originate from a combination of two factors:

(i) Deep neural networks have an inductive bias for "simple" functions that only change significantly upon large-scale changes of the input [164, 245–247]. This is desirable for machine-learning applications like image recognition: the classification of an image should not normally depend on changing a single pixel. However, the same bias is detrimental for learning a nontrivial quantum phase structure.

(ii) The ansatz would typically be initialised with a ferromagnetic sign structure (§4.3.2). Approaching the antiferromagnetic ground state would require setting a large number of phases $\Phi(\boldsymbol{\sigma})$ from 0 to π mod 2π, which may prove difficult for ansätze as flexible as deep neural networks, even if these $\boldsymbol{\sigma}$ follow a regular pattern.

Shallow, fully connected networks may in fact be more successful because Marshall-type sign rules are simple functions of linear combinations of the σ_i^z, to which they have a much stronger bias than deep networks. A case in point is RBMs: since

$$\cosh(x + i\pi) = -\cosh x = e^{i\pi} \cosh x, \tag{4.22a}$$

$$\cosh(x \pm i\pi/2) = \pm i \sinh x = e^{\pm i\pi/2} \sinh x, \tag{4.22b}$$

a single hidden spin coupled to the visible ones analogously to the Huse–Elser construction (4.11) is able to recover the MSR. It is clear, however, that the same rigidity of RBMs is detrimental when it comes to solving the strong sign problem.

In fact, our phase ansatz (4.17) can be thought of as a way to make the same Huse–Elser phase prescription translation-invariant by adding a number of such phases, which all generate the same MSR, as phasors. However, the same may also cap the representational power of (4.17) at that of the Huse–Elser ansatz. Interestingly, making the underlying network deeper in our ansatz leads to similar instabilities to those seen in §4.4.3: as the phase of each phasor is now the output of a deep neural network, they may exhibit the problematic bias for "simple" functions described above.

[†]These architectures are different from the ones normally used in VMC in that they represent the probability that $\operatorname{sgn}[\Psi(\boldsymbol{\sigma})]$ be ± 1. This representation helps avoid the second issue discussed in this section, but it does not translate naturally to any quantity in VMC. However, the continuous tuning of the sign is reminiscent of the ansätze discussed in the previous section.

4.5.2 Strategies to solve the sign problem

As discussed above, the (strong) sign problem is a consequence of variational ansätze treating $\Psi = 0$ as special, which hampers their ability to change sign continuously during imaginary-time evolution under a sign-problematic Hamiltonian. This special treatment is explicit in our ansatz (4.16), as the output of the phase network (4.17) must change by π to change the sign of the wave function; nevertheless, other ansätze that represent the wave function as the exponential of a neural network suffer from the same problem, as the network output still has to jump by $i\pi$.

The key to solving this problem appears to be designing neural-network ansätze that allow switching between positive and negative outputs easily. One may go about doing so in a variety of ways:

- The most straightforward way around the sign problem would be taking a deep neural network to represent the wave function directly. Typical neural networks are, however, additively extensive, unlike wave functions and probability distributions, which are multiplicative (§4.2.2). In practical terms, the dynamic range of wave-function amplitudes (i.e., the ratio of the largest and smallest amplitudes) grows exponentially with system size.[†] I found that neural networks with standard (e.g., ReLU) activation layers fail to represent this dynamic range, even with 6–10 convolutional layers:[‡] VMC optimisation converged to networks whose outputs appeared to span 3–5 decades, compared to the 15–20 decades achieved with the exponentiated network of §4.4.1. Accordingly, the optimisation was extremely noisy and converged to energies far above the ground state.

- This issue, together with the sign problem, is avoided by tensor networks as their outputs are linear combinations of products of terms associated with each lattice site. To achieve the same effect in neural networks, note that a *tree tensor network (TTN)* can be thought of a feed-forward neural network in which activation layers are replaced by pairwise-product pooling. This can be generalised to *polynomial neural networks (PNNs),* where activation layers multiply entries or raise them to integer powers. Such networks have been among the first proposed in the literature [200]; they remove the stringent entanglement limitations of TTNs [166] and, in line with my expectations, improve the available dynamical range over ReLU

[†]This problem is partially mitigated by importance-sampling Monte Carlo, which disregards the lowest amplitudes. However, there are many more low-amplitude basis states than high-amplitude ones, so a very broad dynamic range is still necessary to capture the wave function.

[‡]This issue is reminiscent of the well-known difficulty of representing multiplication using neural networks in a way that extrapolates well to either larger or smaller inputs than those used in training [248].

activations. However, I found that optimising deep PNNs is highly unstable, which may limit their usefulness or require more sophisticated network design.

- In systems like the square-lattice J_1–J_2 model, spin-liquid and other exotic phases are borne out of a competition between incompatible ordered phases that can all be captured in terms of Marshall sign rules (cf. Fig. 4.10). A randomly initialised deep network is not expected to recover any of these rules due to their preference for "simple" functions [246], requiring explicit biasing towards them to stabilise VMC optimisation [205]. I am currently designing an ansatz with several components that are biased according to all the competing sign rules in such a way that their relative importance in different parts of the Hilbert space can be learned without developing frequent sign changes in the neural-network outputs. Preliminary results on the square lattice are promising; finding ways to generalise this idea to geometrically frustrated (e.g., triangular or kagome) lattices is, however, a challenging open question.

- The success of the plain Huse–Elser ansatz (4.5) on unfrustrated systems suggests that the "coarse structure" of amplitudes (that is, an envelope that captures most of the dynamical range of the wave function) can be captured using a very simple exponentiated component. This motivates a neural network architecture of the form

$$\Psi(\boldsymbol{\sigma}) = e^{f(\sigma)} g(\boldsymbol{\sigma}), \tag{4.23}$$

  where f is responsible for representing the "coarse structure," while g represents the sign structure as well as the "fine structure" of amplitudes that is due to interference effects. The latter is analogous to the "soft sign" learned in the supervised protocol of Ref. 198: I believe that performance of this protocol is hampered in frustrated phases by states whose amplitudes is reduced from its envelope value by destructive interferences, naturally captured by $g(\boldsymbol{\sigma}) \approx 0$.

- Recent literature [243] suggests that the stability of VMC for a generic network architecture is improved greatly by pretraining the phase structure, similar to §4.3.2. Namely, one starts by training the phase content $\Psi/|\Psi|$ of an otherwise arbitrary architecture $\Psi(\boldsymbol{\sigma})$, which allows the phases to develop antiferromagnetic character without disturbing the amplitudes $|\Psi|$. Using this protocol on a deep convolutional ansatz similar to (4.21) yielded state-of-the-art energies for spin-liquid models on both square and triangular lattices. It will be interesting to see how this approach performs on more challenging systems and whether it can be extended to the real-valued ansätze described above.

4.5.3 Recurrent neural-network states

Another natural strategy, exploiting recent advances in machine learning, is to use more complex and sophisticated neural-network architectures as NQS ansätze. Cases in point are *recurrent neural networks (RNN),* which represent multivariate probability distributions $p(\sigma_1, \ldots, \sigma_n)$ as the chain of conditional probabilities

$$p(\sigma_1, \ldots, \sigma_n) = p(\sigma_1)p(\sigma_2|\sigma_1)\cdots p(\sigma_n|\sigma_1, \ldots, \sigma_{n-1}), \tag{4.24}$$

and sample each conditional distribution by iterating the same neural network and using the output of the ith iteration to inform the next one about $\sigma_1, \ldots, \sigma_i$ [224, 226]. As each conditional distribution is straightforward to sample, RNNs are able to generate perfectly uncorrelated samples of the overall distribution $p(\boldsymbol{\sigma})$. This is a major improvement over Markov-chain Monte Carlo, especially if $p(\boldsymbol{\sigma})$ has several high-probability regions, separated by prohibitively low probabilities, as is often the case for ordered phases that spontaneously break symmetries [210]. The chain structure (4.24) of RNNs makes them particularly well-suited to one-dimensional and time-series problems [211, 224], but adaptations using 2D convolutional networks have been devised for image recognition [210, 249].

RNNs can be used as efficient NQS ansätze [210–212], provided $p(\boldsymbol{\sigma})$ represents the quantum probability distribution $\left|\langle\boldsymbol{\sigma}|\Psi\rangle\right|^2/\langle\Psi|\Psi\rangle$; that is, a normalised RNN ansatz would take the form

$$\Psi(\boldsymbol{\sigma}) = \sqrt{p_{\mathrm{RNN}}(\boldsymbol{\sigma})}\, e^{i\Phi(\sigma)}. \tag{4.25}$$

Direct sampling of the underlying probability distribution p_{RNN} speeds up VMC optimisation, allowing extremely deep networks to be used [210]. As long as the Hamiltonian has no sign problem, the expressive power of these networks leads to highly accurate energy estimates even for large systems [210, 212].

However, the necessary separation of amplitudes and phases in (4.25) brings back the "strong sign problem" of §4.5.1, limiting the usefulness of RNN wave functions for frustrated systems. Indeed, the RNN ansatz used by Ref. 211 attains substantially less accurate variational energies in the frustrated phase of the 1D J_1–J_2 model, and even at the unfrustrated point $J_2 = 0$, imposing the corresponding MSR leads to significant improvements.

4.5.4 Combining NQS and mean-field wave functions

Fermionic mean-field wave functions have been highly successful in modelling the ground states of both interacting fermion systems and (after Gutzwiller projection) antiferromagnets (§4.1.1), including tackling much of the sign problem. It is, therefore, natural to try and enhance the representational power of these states by introducing neural networks. This can be done in a number of different ways:

- Multiply the mean-field wave function with a "neural-network Jastrow factor" [cf. (4.4)]. This is computationally simple for shallow-network ansätze, such as RBMs [216] or Gaussian processes [213], and leads to significant improvements in variational energies. However, owing to the "strong sign problem" of the NQS factors, they tend not to improve the sign structure of the mean-field ansatz, limiting these approaches to the accuracy of fixed-node DMC.

- The "single-particle orbitals" that enter the mean-field Slater determinant can be made dependent on the positions of the other particles:

$$\Psi_{\text{backflow}}(\vec{r}_1,\ldots,\vec{r}_n) = \begin{vmatrix} \psi_1(\vec{r}_1;\vec{r}_2,\ldots,\vec{r}_n) & \cdots & \psi_n(\vec{r}_1;\vec{r}_2,\ldots,\vec{r}_n) \\ \vdots & \ddots & \vdots \\ \psi_1(\vec{r}_n;\vec{r}_1,\ldots,\vec{r}_{n-1}) & \cdots & \psi_n(\vec{r}_n;\vec{r}_1,\ldots,\vec{r}_{n-1}) \end{vmatrix}. \tag{4.26}$$

  If the orbitals ψ_i are invariant under permutations of its variables after the semicolon, Ψ_{backflow} remains a valid, antisymmetric, fermionic wave function. This *backflow* approach dates back to Feynman [250]; neural networks can be used as additive [251] or multiplicative [208] corrections to the single-particle orbitals. Recently, several sophisticated deep-network ansätze based on (4.26) have been designed and used to obtain molecular ground-state energies to within chemical accuracy [207, 208].

For antiferromagnetic systems, a particularly promising approach was proposed in Ref. 175, which combined pair-product states[†] with an RBM many-body correlator:

$$|\Psi'\rangle = \prod_k \cosh\left(b_k + \sum_i w_{ik}\hat{\sigma}_i^z\right) \times \hat{P}_{\text{Gutzw.}}\left(\sum_{i,j} f_{ij} c_{i\uparrow}^\dagger c_{j\downarrow}^\dagger\right)^{N_{\text{site}}/2} |\text{VAC}\rangle, \tag{4.27}$$

where the w and the b are the weights and biases of the RBM, $\hat{P}_{\text{Gutzw.}}$ is the Gutzwiller projection operator, f_{ij} are variational parameters of the pair-product state, $c_{i\uparrow\downarrow}^\dagger$ creates

[†]These are analogous to the ground states of BCS mean-field Hamiltonians with hopping and pairing terms [38].

an up- or down-spin fermion on site i, N_{site} is the number of lattice sites, and $|\text{VAC}\rangle$ is the fermionic vacuum state. The RBM sign problem is resolved by making $|\Psi'\rangle$ translation invariant using (4.21), cf. §4.5.1. The resulting wave function achieves excellent variational energies in the spin-liquid phase of the square-lattice J_1–J_2 model, both for the ground state and excited states.

One may wonder, however, how much these approaches make good on the original promise of NQS, the ability to obtain accurate ground states without much physical insight. Exchange antisymmetry is a key property of fermion wave functions: using Slater determinants to enforce it is not very different from using convolutional networks for translation-invariant problems.[†] By contrast, fermionic mean-field wave functions are relevant for spin systems due to the connection between Heisenberg and Hubbard models (§1.2, §4.1.1): as such, they are far from unbiased even for Heisenberg models, and may do rather poorly for anisotropic systems (e.g., quantum spin ice), for which similar parton constructions are not readily available.

[†]They are, however, not indispensable. Other approaches, closer in vein to Laughlin wave functions, can also be used to enforce antisymmetry, with similarly accurate results [206].

5

Monopoles and antiferromagnetic domain control in spin-ice iridates

In spite of extensive experimental and theoretical work on spin ices, a reliable indicator of the density of their emergent monopoles is yet to be found. In this chapter, I discuss magnetisation and magnetoresistance experiments on the spin-ice material $Ho_2Ir_2O_7$, together with our theoretical interpretation of them, supported by Monte-Carlo simulations. Our results show that the magnetoresistance of $Ho_2Ir_2O_7$ is highly sensitive to the density of monopoles and could be used to quantitatively measure it in experiments, providing a new, powerful tool to study spin-ice systems.

Furthermore, for certain orientations, external magnetic fields couple strongly to the antiferromagnetically ordered iridium ions via the large holmium moments. The result is the ability to manipulate antiferromagnetic domains with magnetic fields, which is evidenced by a strong hysteresis that appears simultaneously in both the magnetisation and the magnetoresistance response. The precise control of antiferromagnetic domain walls is a key goal for the design of next-generation spintronic devices [252]. Our results show a way forward to how this target may be achieved using externally applied magnetic fields and the peculiar interplay between frustrated localised moments and itinerant antiferromagnetism in rare-earth oxide materials.

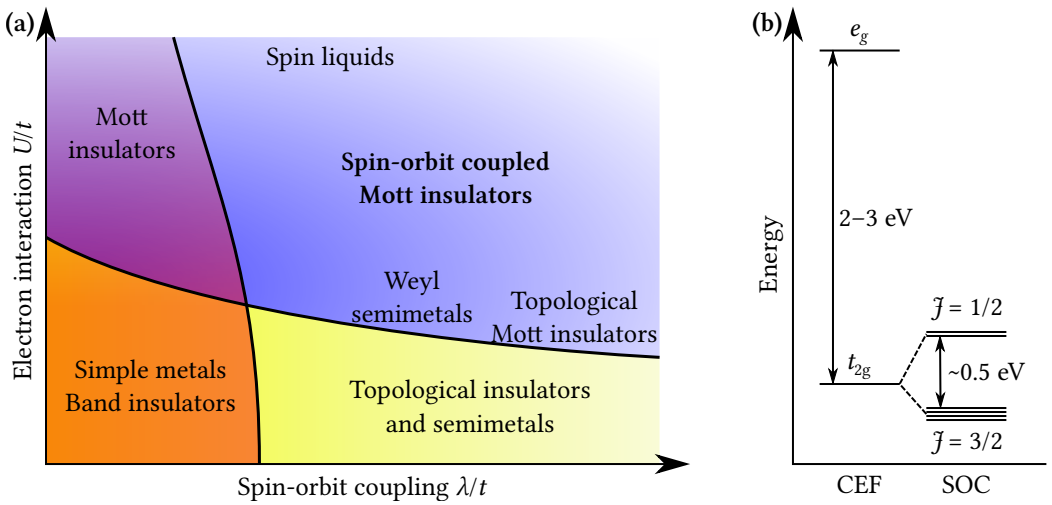

Figure 5.1. **(a)** Schematic phase diagram of the extended Hubbard model (5.1). In 5*d* materials, both the interaction strength U and the spin-orbit coupling λ are large, leading to chiral Mott insulators [253].
(b) Iridium 5*d* energy levels in an octahedral CEF environment [254]. In the Ir^{4+} ion, five *d*-electrons fill all four $J = 3/2$ and one $J = 1/2$ orbitals: the latter becomes an effective spin-1/2 degree of freedom with strong spin–orbital correlations.

5.1 Introduction: magnetism in iridium compunds

There's no limit to how complicated things can get, on account of one thing always leading to another.

— E. B. White

Iridium and other 5*d* elements play a prominent role in many recent developments of strongly correlated electron physics. They have a much stronger spin–orbit coupling (SOC) than light 3*d* elements: this is a key ingredient for engineering topologically non-trivial band structures in heavy *p*-electron materials [255]. In the more tightly localised *d*-orbitals, these effects compete with those due to electron–electron interactions, captured by the extended Hubbard model

$$H = \sum_{\substack{ij,\alpha\beta \\ \sigma=\uparrow,\downarrow}} (t_{ij,\alpha\beta} c^\dagger_{i\alpha\sigma} c_{j\beta\sigma} + \text{H.c.}) + \lambda \sum_i \vec{L}_i \cdot \vec{S}_i + U \sum_{i\alpha} n_{i\alpha\uparrow} n_{i\alpha\downarrow}; \tag{5.1}$$

in 5*d* systems, both U and λ are comparable to t, leading to a variety of spin–orbit coupled Mott-insulator phases [253] [Fig. 5.1(a)]. Finally, correlations between spin and orbital quantum numbers due to SOC, together with the strong orbital anisotropy of *d*-electron hopping, result in highly anisotropic spin interactions, allowing exotic spin-

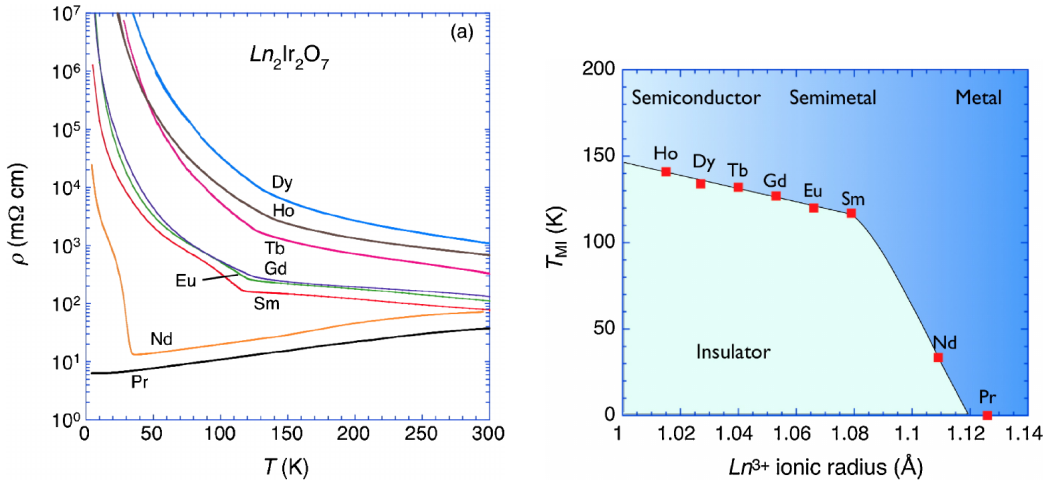

Figure 5.2. Left: resistivity as a function of temperature in several rare-earth pyrochlore iridates. All but $Pr_2Ir_2O_7$ show a metal–insulator transition between 30 K and 150 K. For Nd, Sm, and Eu, there is a sudden upturn in resistivity at the transition, while Gd, Tb, Ho, and Dy only show a modest shoulder.
Right: metal–insulator transition point as a function of rare-earth ionic radius. The sudden reduction of T_{MI} suggests a quantum (semi)metal–insulator transition.
Figures taken from Ref. 258.

liquid phases to be realised. For example, Ir^{4+} and Ru^{3+} cations become effective spin-1/2 degrees of freedom in an octahedral CEF environment, with $\sigma^{x,y,z}$ spin components correlated with the $d_{yz,xz,xy}$ orbitals, respectively [Fig. 5.1(b)]. The same octahedral environments also result in highly anisotropic couplings between the effective spins, leading to Kitaev-type physics in α-$RuCl_3$ and $(Li, Na)_2IrO_3$ [256, 257].

5.1.1 Pyrochlore iridates

This chapter focuses on rare-earth iridates of empirical formula $R_2Ir_2O_7$, where R is a 3+ rare earth cation. Apart from $La_2Ir_2O_7$, these all form a pyrochlore lattice with R^{3+} ions on the A sites and Ir^{4+} ions on the B sites (Fig. 1.5) [41]. The latter have octahedral oxygen coordination similar to that in $(Li, Na)_2IrO_3$ [Fig. 1.6(e)]; accordingly, the crystal-field and spin–orbit splittings give rise to an effective spin-1/2 moment [Fig. 5.1(b)]. Unlike the rare-earth f-moments, these effective spins remain approximately isotropic. Accordingly, the exchange interactions between nearest neighbours are dominated by isotropic Heisenberg exchange; as the iridium d-orbitals are much less strongly localised than rare-earth f-orbitals, these exchange couplings are much stronger (on the order of 500 K). The pure Heisenberg antiferromagnet on the pyrochlore lattice is a spin liquid [92, 259, 260]; however, an additional direct Dzyaloshinskii–Moriya interaction brings about

an ordering transition within the ground-state manifold, where the iridium moments assume an all-in-all-out (AIAO) antiferromagnetic order [261]. As a result, all rare-earth iridates apart from $Pr_2Ir_2O_7$ undergo an AIAO ordering transition at 30–150 K (Fig. 5.2) [258].

Compared to the strongly insulating titanates, pyrochlore iridates have fairly low resistivity at high temperatures, mostly due to iridium conduction electrons. The magnetic ordering transition, however, induces a metal–insulator transition, too [258, 262, 263]; the corresponding upturn in resistivity is rather pronounced for the lighter rare-earth elements Nd, Sm, and Eu,[†] but much less so for Gd, Tb, Dy, and Ho (Fig. 5.2) [258]. Their low-temperature electronic structure is an exciting playground for topological and strongly correlated electronic physics: the band structure exhibits a large number of Weyl nodes [264], with corresponding Fermi arcs forming on grain boundaries, as well as domain walls [265]. Recently, $Nd_2Ir_2O_7$ has also been argued to show non-Fermi-liquid behaviour [266].

5.1.2 Fragmented spin ice in $Ho_2Ir_2O_7$

In our work, we studied the low-temperature magnetism of the pyrochlore iridate $Ho_2Ir_2O_7$. Due to the similar size of Ti^{4+} and Ir^{4+} ions, the exchange pathways and local crystal fields acting on the Ho^{3+} ions are similar to those in $Ho_2Ti_2O_7$; accordingly, interactions between the latter are described well by the titanate dipolar-spin-ice Hamiltonian (1.29). In addition, the oxygen ions mediate an exchange coupling between the holmium and iridium ions, of a similar strength as the spin-ice couplings [69]. However, the temperatures where either of these interactions is relevant (below 10 K) are far below the metal–insulator transition point (around 150 K): at these temperatures, the iridium moments form perfect AIAO domains with sparse domain walls and negligibly few thermal defects. By symmetry, the exchange couplings due to this iridium configuration would drive the holmium moments to a similar AIAO order [Fig. 5.3(b)]. Therefore, the holmium spins can be modelled as a classical spin ice with an additional effective field, pointing along the local ⟨111⟩ axis, due to the iridium environment [69]:

$$H = \frac{J}{3}\sum_{\langle ij\rangle}\sigma_i\sigma_j + D\ell^3\sum_{ij}\sigma_i\sigma_j\left[\frac{\hat{e}_i\cdot\hat{e}_j}{r_{ij}^3} - \frac{3(\hat{e}_i\cdot\vec{r}_{ij})(\hat{e}_j\cdot\vec{r}_{ij})}{r_{ij}^5}\right] \pm h_{\mathrm{loc}}\sum_i\sigma_i. \qquad (5.2)$$

The terms in (5.2) correspond to holmium–holmium nearest-neighbour and dipolar interactions and holmium–iridium coupling, respectively; the sign of the last term depends on which of the two possible AIAO domains are realised by the iridium spins.

[†]Even in these cases, the metal–insulator transition can be suppressed by substituting iridium on a small fraction of the rare-earth sites, indicating the proximity of the insulating phase to a metallic one [262].

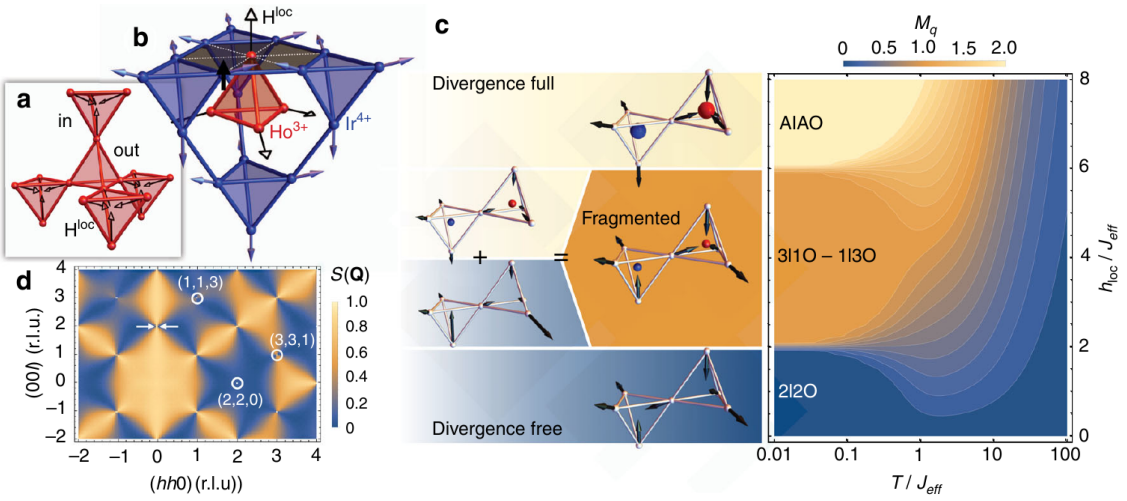

Figure 5.3. **(b)** Arrangement of Ho^{3+} (red) and Ir^{4+} (blue) ions in $Ho_2Ir_2O_7$. Each holmium ion is surrounded by a hexagon of six iridium nearest neighbours. When the iridium lattice orders magnetically in the AIAO phase, as shown by the blue lattice, an effective field along the local ⟨111⟩ directions is felt by the holmium moments in the centre, as shown in **(a)**.
(c) Phase diagram of the nearest-neighbour spin-ice iridate model ($J_{\text{eff}} = -J/3$). At small h_{loc}, the two-in-two-out (2I2O) spin-ice state remains stable (blue); very large fields override the interactions to create an AIAO-ordered state (yellow). For intermediate h_{loc}, the ground state is a 3I1O–1I3O monopole crystal (orange), where the $\sigma = \pm 1$ Ising moments fragment into AIAO ordered $\pm(1/2, 1/2, 1/2, 1/2)$ and divergence-free $\pm(1/2, 1/2, 1/2, -3/2)$ components. These phases remain distinct up to $T \sim J_{\text{eff}}$; high temperatures wash the AIAO order parameter M_q out.
(d) Spin correlation function $S(\vec{q})$ calculated in the fragmented phase. It shows a diffuse component with pinch points (arrows) together with magnetic Bragg peaks (circled) characteristic of the AIAO-ordered state. The intensity of these (proportional to the ordered moment squared) is a quarter of that expected for full AIAO order.
Figure taken from and caption based on Ref. 69.

The low-temperature behaviour of $Ho_2Ir_2O_7$ was first studied in Ref. 69. Using powder neutron scattering, they found that holmium spins[†] develop an ordered moment consistent with AIAO ordering at low temperature; the magnitude of this moment is, however, only half that corresponding to full ordering. This can be understood in terms of *moment fragmentation:* the ground state is a monopole crystal with alternating three-in-one-out (3I1O) and one-in-three-out (1I3O) tetrahedra, which can be regarded as the equal-weight superposition of an AIAO state and an arrangement of +1 and −3 "Ising

[†]The magnetic moment of the iridium spins orders with the same wave vectors; however, it is orders of magnitude smaller than that of the holmiums [258, 267], so it can be safely neglected in neutron-scattering and magnetisation experiments.

spins" that obeys a divergence-free constraint. The former yields the observed AIAO-ordered moment, while the latter has an exponential number of ground states, resulting in extensive zero-point entropy, pinch points in neutron scattering [Fig. 5.3(d)], and strong paramagnetic response. Using a nearest-neighbour-only model [that is, Eq. (5.2) without the D term], Ref. 69 demonstrated analytically that this monopole crystal is stable for a finite range of h_{loc}: this range includes both $Ho_2Ir_2O_7$ and the analogous "spin-ice iridate" $Dy_2Ir_2O_7$ [268].

5.2 Experiments on $Ho_2Ir_2O_7$

The true method of knowledge is experiment.

— William Blake

> **Note:** the experiments discussed in this section and §5.4 were performed without my involvement. I report them in detail because they form the immediate background of my theoretical work in §5.3 and §5.4.

Small single crystals of $Ho_2Ir_2O_7$ were grown by D. Prabhakaran's group in Oxford. First, polycrystalline $Ho_2Ir_2O_7$ powder was prepared using high purity (>99.99%) Ho_2O_3 and IrO_2 chemicals in molar ratio $1 : 1.05$ (excess IrO_2 was added to compensate for evaporation loss). The powders were thoroughly mixed along with 0.1 g KF per 5 g powder inside an argon glove box and pressed into pellets of diameter 15 mm. These were sealed inside a platinum crucible and sintered at 1100°C for 100 hours. The resulting phase-pure[†] $Ho_2Ir_2O_7$ powder was mixed with KF flux in ratio $1 : 200$ and packed into a platinum crucible with a tightly fitted lid [269]. The crucible was placed inside a chamber furnace and heated to 1050°C and, after holding for 10 hours, it was cooled down to 850°C at rate 1°C/h, and finally to room temperature at 60°C/h. Octahedral-shaped single crystals of size up to 1 mm^3 were separated after dissolving the flux with hot water; in the experiments that follow, two crystals of size $\approx 0.2 \times 0.2 \times 0.2$ mm^3 and mass ≈ 0.1 mg were used. Magnetisation, susceptibility, and magnetoresistance measurements on these samples were performed by Paul Goddard's group at Warwick.

The magnetisation (M vs. H) and susceptibility (χ vs. T) of a single crystal of $Ho_2Ir_2O_7$ were measured using a Quantum Design MPMS superconducting quantum interference device (SQUID) magnetometer, except for the magnetisation data at different field sweep rates [Fig. 5.7(a)], which were collected using an Oxford Instruments vibrating sample magnetometer (VSM). VSM magnetisation data were calibrated using a multiplicative

[†]Phase purity of the powder and the final single crystals were characterised using PANalytical and Supernova x-ray diffractometers, respectively.

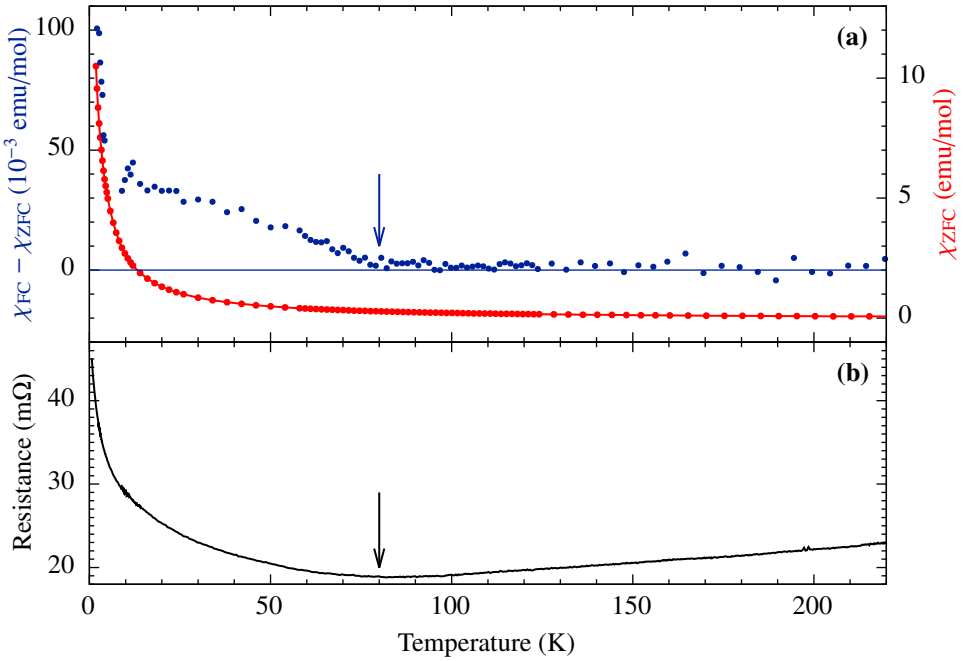

Figure 5.4. **(a)** Susceptibility measurements on $Ho_2Ir_2O_7$ as a function of temperature in zero (ZFC, red dots) and 0.01 T (FC) fields. They are visually indistinguishable at all temperatures, so $\Delta\chi = \chi_{FC} - \chi_{ZFC}$ (blue dots) is plotted instead of χ_{FC}. Above $T_{MIT} \approx 80$ K (arrow), $\Delta\chi \approx 0$, but it becomes finite below, indicating an (antiferromagnetic) ordering transition.
(b) Resistance measurements on $Ho_2Ir_2O_7$ as a function of temperature in zero external field. The resistance has a minimum around 80 K (arrow), rising sharply underneath, indicating a metal–insulator transition coincident with magnetic ordering. The increase in resistivity in the insulating phase is less dramatic than in Ref. 258.
Figure based on Ref. 4.

scale factor derived from SQUID measurement at the same temperature. All magnetometry measurements were performed on the same single crystal of mass 134(5) μg.

Electrical-resistance measurements were made using a four-wire technique with AC current of magnitude 855 μA and frequency around 150 Hz. Magnetic fields were applied using an Oxford Instruments superconducting magnet equipped with a ^{3}He insert. All resistance measurements were performed on the same single crystal of $Ho_2Ir_2O_7$, taken from the same growth batch as the crystal used for the magnetometry measurements.

The iridium AIAO ordering transition was detected by measuring the susceptibility of a $Ho_2Ir_2O_7$ sample both in zero and 0.01 T magnetic fields [Fig. 5.4(a)]: the two curves bifurcate below 80 K due to spontaneous symmetry breaking. The measured resistance has a minimum at the same temperature [Fig. 5.4(b)], which indicates the metal–insulator transition concomitant to magnetic ordering. The transition temperature is substantially lower and the upturn of the resistivity in the insulating phase is much more gentle in

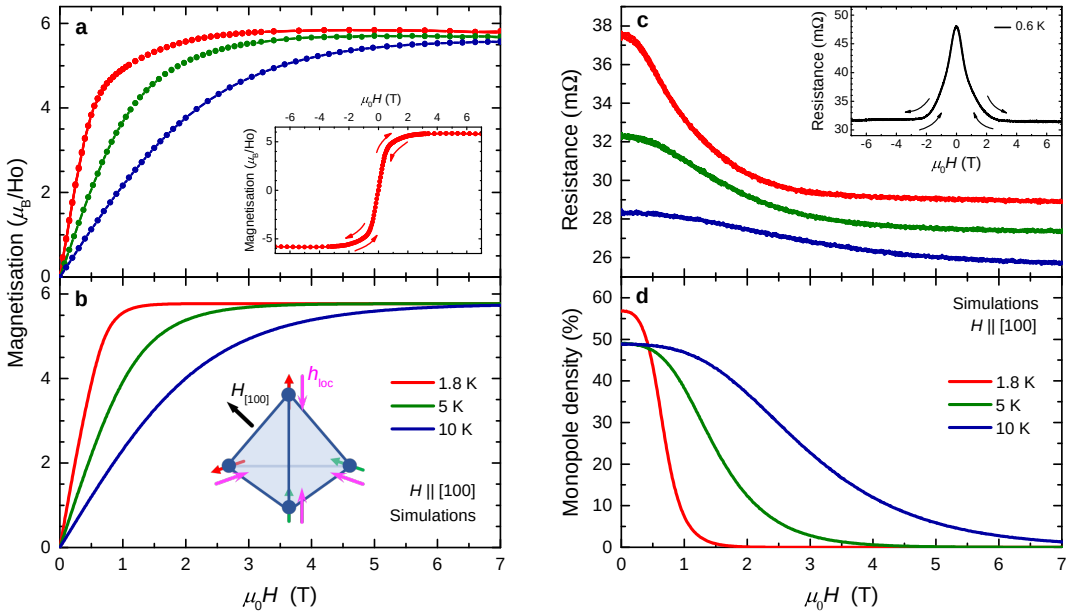

Figure 5.5. $Ho_2Ir_2O_7$ under an external field along the [100] crystallographic axis.
(a) Measurements and **(b)** Monte-Carlo simulations of the magnetisation at various temperatures.
Inset to (a): magnetisation measurement at 1.8 K upon sweeping the field continuously between ±7 T.
Inset to (b): single tetrahedron of the Ho^{3+} sublattice with local fields $\vec{h}_{loc}$ (magenta arrows) due to the ordered iridium moments. Under a large external [100] magnetic field (black arrow), the holmium moments order into a 2120 configuration, with two moments parallel (green arrow) and two antiparallel (red arrow) to $\vec{h}_{loc}$.
(c) Resistance measurements and **(d)** Monte-Carlo simulations of the single-monopole density at various temperatures. The negative magnetoresistance is due to the suppression of magnetic disorder in general, and of magnetic monopoles in particular.
Inset to (c): resistance measurement at 0.55 K upon sweeping the field continuously between ±8 T.
Figure taken from Ref. 4.

our work than in Ref. 258: this might be explained by subtle differences in sample quality [262].

5.2.1 [100] magnetic field

The magnetic field was first applied along the [100] crystallographic direction. The phenomenology for this orientation underpins the richer and more complex behaviour when the field is applied along the [111] direction, discussed in §5.2.2.

The magnetisation rises rapidly under the applied field, reaching a saturation value

$M^{sat}_{[100]} = 5.8(2)\mu_B$/Ho at 1.8 K [Fig. 5.5(a)]. This is consistent with holmium moments of magnitude $\mu_{Ho} = 10.1(4)\mu_B$, as the magnetisation of the ordered 2I2O state expected in the field-polarised limit [inset to Fig. 5.5(b)] is $\mu_{Ho}/\sqrt{3}$ per atom [270, 271] and the iridium magnetism has little qualitative or quantitative effect on the magnetisation response [267]. On sweeping the field continuously between ±7 T [inset to Fig. 5.5(a)], no hysteresis between sweeping up and down is observed.

Figure 5.5(c) shows the electrical resistance of the $Ho_2Ir_2O_7$ sample as a function of an applied [100] magnetic field. From a temperature-dependent starting value, we first observe a large negative magnetoresistance which flattens out at higher fields. The zero-field resistance increases as the temperature is lowered since $Ho_2Ir_2O_7$ is insulating below the metal–insulator transition at 80 K. The negative magnetoresistance is caused by a reduction in scattering off the holmium moments as they order into a regular 2I2O state under the applied magnetic field. Other than this state, there are no (partially) ordered spin configurations that arise at intermediate [100] fields: as a result, the magnetoresistance curve is rather featureless, reaching a constant high-field value as the magnetisation saturates. A continuous sweep of the field between ±7 T yields no observable hysteresis [inset to Fig. 5.5(c)], similar to the magnetisation curves.

As the temperature is increased, the initial negative magnetoresistance broadens out to higher fields. This is to be expected because at higher temperatures, like in any paramagnet, larger fields are required to order the holmium moments. The same is reflected in the slower saturation of magnetisation at higher temperatures [Fig. 5.5(a)].

5.2.2 [111] magnetic field: magnetisation and resistivity hysteresis

Figure 5.6(a) shows that in an external magnetic field along the [111] direction, the magnetisation rises to saturation at $M^{sat}_{[111]} = 5.1(2)\mu_B$/Ho at 1.8 K. This is consistent with $\mu_{Ho} = 10.1(4)\mu_B$, since the expected magnetisation of the field-polarised 3I1O–1I3O monopole crystal (Fig. 5.8) is $\mu_{Ho}/2$ [270, 271]. The dominant source of error in both estimates of μ_{Ho} is that of the sample mass; to cancel this error out, we compare the saturation magnetisation in [100] and [111] fields: $M^{sat}_{[111]}/M^{sat}_{[100]} = 0.872(5)$, in excellent agreement with the expected $\sqrt{3}/2 \approx 0.866$.

A striking observation is the presence of hysteresis in the magnetisation for this orientation. Notably, the hysteresis is closed at zero applied field, which indicates that the field sweeps are slow enough for the holmium moments to remain in equilibrium. This is further confirmed by magnetisation measurements performed at a range of sweep rates between 0.05 and 1 T/min [Fig. 5.7(a)]: the hysteresis loop remains nearly identical for all sweep rates, showing that the hysteresis opens due to a quasi-equilibrium process.

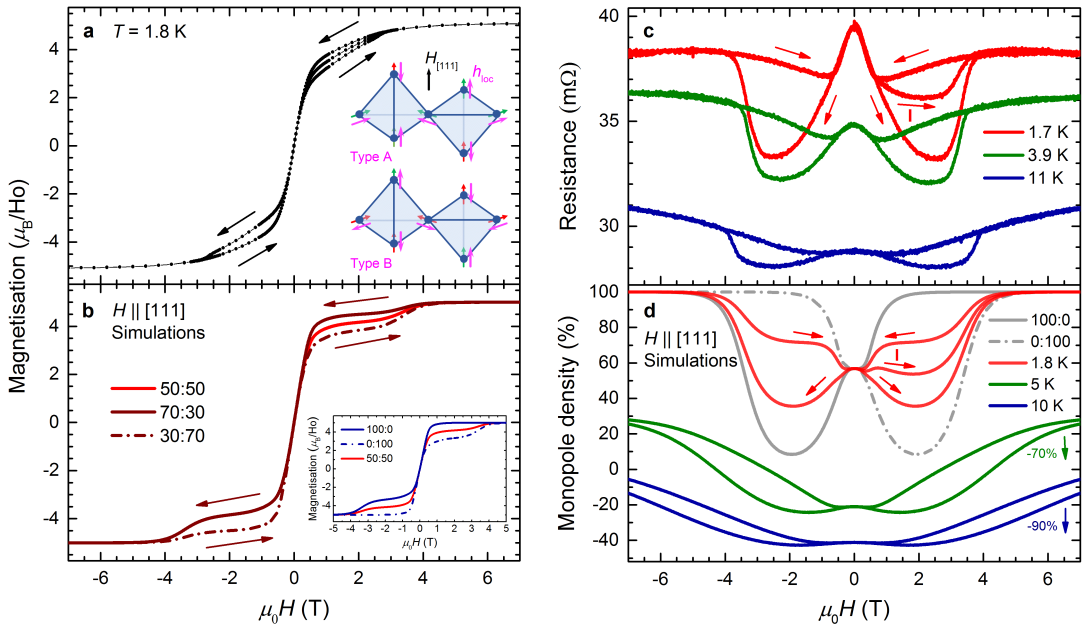

Figure 5.6. $Ho_2Ir_2O_7$ under an external field along the [111] crystallographic axis.
(a) Measurements and **(b)** Monte-Carlo simulations of the magnetisation at various temperatures. Arrows show the direction of each sweep; the additional curve lying between the upsweep and downsweep for $\mu_0 H > 0$ is the initial (virgin) sweep after cooling the sample in zero field.
The 100 : 0 and 0 : 100 curves in the **inset to (b)** show simulations on type-*A* and *B* single domains, respectively. To approximate domain plasticity in **(b)**, I used 70 : 30, 30 : 70, and 50 : 50 weighted averages of these for the up-, down-, and virgin sweeps, respectively.
(c) Resistance measurements and **(d)** Monte-Carlo simulations of the single-monopole density at various temperatures. Arrows show the direction of each sweep, the virgin sweep is marked with "I" (lowest temperature only). Monopole densities in **(d)** on the up-, down-, and virgin sweeps are approximated with 70 : 30, 30 : 70, and 50 : 50 weighted averages of simulations of pure type-*A* and *B* domains. The 5 K (10 K) curves in **(d)** have been shifted down by 40% (60%) for clarity. There is no corresponding offset in **(c)**; the curves are shifted by the temperature dependence of resistivity (cf. Fig. 5.5).
Figure taken from Ref. 4.

The hysteresis opens at a finite value of the applied field: for example, at 1.8 K, the hysteresis becomes noticeable around 0.2 T and then closes at 2.9 T. The first ("virgin") field sweep after cooling from above the metal–insulator transition in zero magnetic field sits between the subsequent downsweeps and upsweeps of the hysteresis loop. Other than this sweep, the positive and negative quadrants of the magnetisation loop are symmetric within experimental error.

To gain further insight into the nature of the magnetisation hysteresis, its width (that is, the difference between downsweep and upsweep magnetisation) was measured

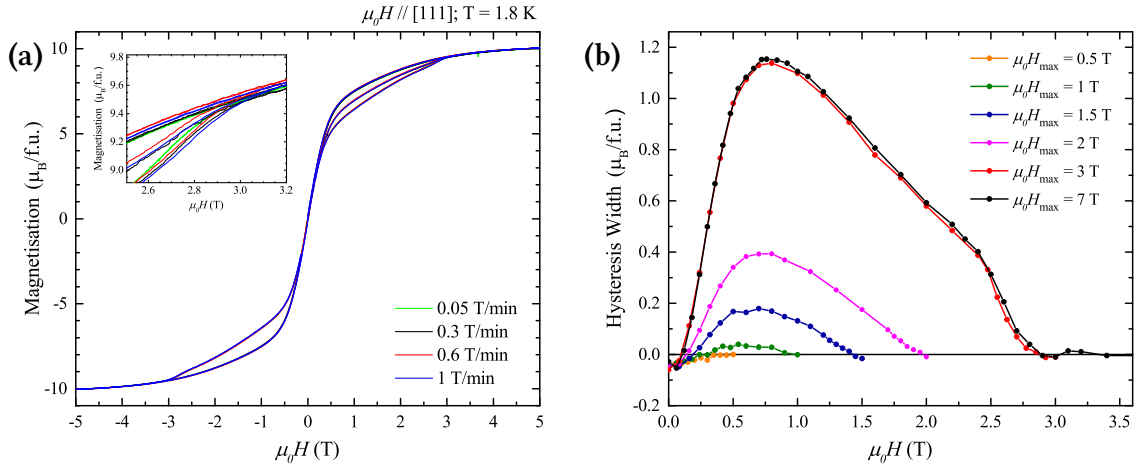

Figure 5.7. **(a)** Magnetisation measurements on $Ho_2Ir_2O_7$ in an external [111] field swept at different rates. The inset zooms in on the region where the hysteresis bubble closes. There is negligible difference between the curves, underlining that our hysteresis is not dynamic in origin.
(b) Hysteresis width in $Ho_2Ir_2O_7$ magnetisation measurements under an applied [111] magnetic field at 1.8 K, defined as downsweep minus upsweep magnetisation. For each curve, the hysteresis loop is performed to a different maximum field. Hardly any hysteresis opens up to about 1 T maximum field, and even at 2 T, its width is much smaller than for higher fields. The shape of the bubble no longer changes for fields above 3 T, where it also closes.
Figures taken from Ref. 4.

for sweeps reaching up to a variety of maximum fields [Fig. 5.7(b)]. Up to about 1 T, no observable hysteresis opens, and even at $\mu_0 H_{\text{max}} = 2$ T, the hysteresis width is only a third of that for the highest fields. By contrast, sweeping above the closing point of 2.9 T does not affect the hysteresis loop. This implies that the hysteresis is due to a quasi-equilibrium (e.g., plastic) process that occurs at external fields between 1 T and 3 T.

Figure 5.6(c) shows the magnetoresistance in [111] magnetic fields, which, like the magnetisation, is highly hysteretic. In particular, the resistance during upsweeps is reduced substantially (by up to 15% at 1.8 K) compared to the downsweeps between about 0.5 T and 3.5 T, in the range where the magnetisation is also hysteretic. The virgin curve (I), shown only for the 1.8 K measurement, lies between the subsequent downsweeps and upsweeps. Similar to the [100] case, the magnetoresistance curves get wider with increasing temperature (magnetisation curves are only shown at 1.8 K to avoid clutter): this is again due to the broadening of paramagnetic response at higher temperatures and the correspondingly greater difficulty of ordering the holmium moments into a 3I10–1130 monopole crystal.

5.3 Theoretical interpretation

> *Nothing has such power to broaden the mind as the ability to investigate systematically and truly all that comes under thy observation in life.*
>
> — Marcus Aurelius, *Meditations iii.11*

In this section, I focus on two key features of the experimental results above:

- $Ho_2Ir_2O_7$ shows a substantial, sweep-rate-independent magnetisation hysteresis in a [111] magnetic field;
- it exhibits a large hysteresis in magnetoresistance too, which can be separated from the saturation of magnetisation in a [110] field (§5.4).

Neither phenomena have previously been observed in pyrochlore iridates and, accordingly, no theoretical explanation for them is available.

The stability of the hysteresis in [111] fields suggests a non-dynamical origin, with relaxation times on the order of hours at least. The dynamics of holmium magnetism is much faster at these temperatures [272]; furthermore, a dynamical hysteresis would be open at zero field, which is not observed here. By contrast, the AIAO order of iridium spins naturally breaks up into several domains of the two symmetry-broken ordered states (Fig. 5.8): similar to standard ferromagnetic domains [273], these may deform under forces applied to domain walls, causing hysteretic changes that also affect holmium magnetism and are stabilised by domain-wall pinning due to disorder, dislocations, etc., accounting for its quasi-equilibrium nature. No direct force, however, is exerted on the antiferromagnetic domains by an external field, as they have zero net magnetisation. By contrast, I demonstrate in §5.3.2 that the field-polarised 3110-1130 monopole crystal of holmium moments induced by [111] fields does give rise to such a force, thereby providing a coupling between the external field and the iridium AIAO order *mediated by the holmium spin ice.*

I argue furthermore that the hysteresis observed in magnetoresistance is due to changes in scattering off the magnetic texture of the holmium ions, particularly the magnetic monopoles in the spin-ice structure. I show that most of the magnetic scattering of conduction-electron spins off the localised holmium moments can be accounted for in terms of spin-ice monopoles described by the dumbbell model [62]. In addition, spin-ice materials like $Dy_2Ti_2O_7$ were shown to be multiferroic [274–276]; in particular, their magnetic monopoles cause charge rearrangement that results in localised electric dipoles [277], which in turn act as scatterers for conduction electrons. I found that these

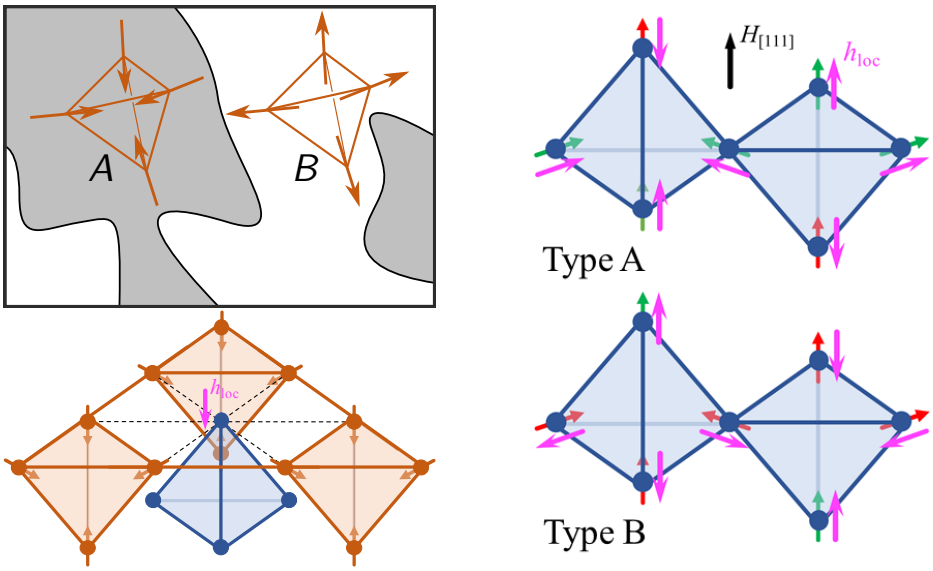

Figure 5.8. **Left**: There are two inequivalent iridium AIAO orders in which the spins point into either "up" or "down" tetrahedra. Upon cooling through the metal–insulator transition, these naturally form a multidomain structure. The direction of the effective Ho–Ir coupling field $\vec{h}_{\mathrm{loc}}$ (magenta arrow) is determined by the orientation of neighbouring iridium spins, so it is different in the two domains.
Right: In a large [111] magnetic field (black), the spins order in a 3in–1in monopole crystal. In type-*A* domains, three polarised spins per tetrahedron (green) are aligned with h_{loc}, compared to just one in type-*B* domains. This induces an energetic splitting between the two domain types that acts as an effective force on domain walls. Upon reversing the external field, the direction of this force is reversed, too.
Figures taken from Ref. 4.

two effects are comparable in strength, and may realistically account for the observed resistance-hysteresis width.

These theoretical arguments are corroborated by classical Monte-Carlo simulations of the holmium-only effective model (5.2), details of which are given in Appendix D. In particular, I demonstrate that the magnetic-field dependence of monopole density is drastically different in type-*A* and *B* domains, leading to a similar hysteresis loop as seen in the resistance measurements upon cycling the external [111] field.

5.3.1 Magnetisation hysteresis in [111] fields

Iridium domains in $Ho_2Ir_2O_7$ are coupled to the spin ice made out of the large holmium moments, which shows a strong response to external fields. In a [111] field in particular, the moments polarise towards a 3in–1in ordered monopole crystal, with alter-

nating kagome and triangular layers pointing in and out of "up" tetrahedra, respectively (Fig. 5.8). In a strong enough field, all spins become fully polarised, regardless of the direction of the exchange field h_{loc}; however, the field direction shown in Fig. 5.8 favours type-A domains since h_{loc} aligns with it for three spins per tetrahedron, compared to just one in a type-B domain.

If the crystal were to consist of a single, fixed domain, this effect would lead to asymmetric magnetisation, as shown in the inset to Fig. 5.6(b), where the 100 : 0 curve shows magnetisation simulations for single type-A domain. In this case, the external [111] field and h_{loc} favour the same monopole crystal, so the holmium moments saturate to it rapidly. A $[\overline{111}]$ field ($H < 0$ in the figure) applied to the same domain competes with the local exchange field, so it must be swept to a higher value to fully align the moments. This realignment happens via an intermediate regime in which monopoles are depleted, resulting in a magnetisation plateau prior to saturation. The 0 : 100 curve shows the converse behaviour for a single type-B domain. The average of these two lines is shown in the 50 : 50 curve, which simulates the response of a crystal that contains a fixed and equal mixture of both iridium domains: it is symmetric but not hysteretic. It is clear that none of these scenarios account for the observed symmetric and hysteretic behaviour alone: plastic deformation of domain boundaries must be considered.

The wide "hysteresis loop" shown in the inset is, however, inconsistent with the experimental hysteresis width. Comparing the widest points of the experimental and numerical hysteresis curves in Fig. 5.6 suggests that the iridium moments are not swept to single type-A or B domains, but their ratio is limited to about 70 : 30 at strong fields. Weighted averages corresponding to such mixtures, and the virgin sweep with an equal split of domains, are shown in Fig. 5.6(b).

5.3.2 Antiferromagnetic domain control using [111] fields

Usually, antiferromagnetic domains cannot be manipulated using external magnetic fields along any direction, as their net magnetisation is zero.[†] In $Ho_2Ir_2O_7$, however, the asymmetry between domains in a [111] field provides a driving force for domain-wall movement. In the fully field-polarised limit, three and one spins point in the direction favoured by h_{loc} in type-A and B domains, respectively; therefore, the Ho–Ir interaction energy of the two domains becomes $\pm 2h_{\mathrm{loc}}$ per tetrahedron: the lower energy of type-A domains generates an effective pressure on the domain walls that causes the same domains to grow. Upon reversing the field direction, type-B domains become favoured

[†] In large fields, spin canting and similar effects may bring about some magnetisation; these higher-order effects are, however, rather weak [278].

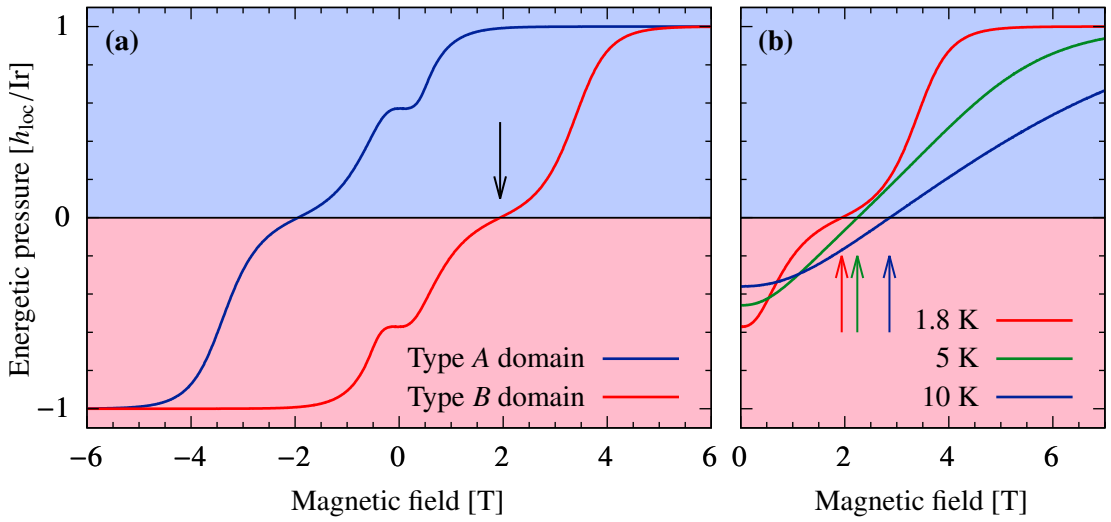

Figure 5.9. **(a)** Energetic pressure $2h_{loc}\langle\sigma\rangle$ due to holmium moments in equilibrium inside both type-A and B iridium domains as a function of external [111] fields at 1.8 K, obtained from Monte-Carlo simulations. For small fields, h_{loc} stabilises both domain types; large fields polarise both domains to the same 3110–1130 monopole crystal. The energetic pressure switches sign around 1.9 T (black arrow), so plastic domain-wall movement speeds up around this field.
(b) Energetic pressure in type-B domains at three different temperatures. Its field dependence broadens significantly at higher temperatures; by contrast, the point where the pressure changes sign shifts up only slightly, matching the evolution of the field at which the experimental hysteresis loops close.
The background shading indicates which domain type the energetic pressure favours: A (light blue) or B (light red). Figures based on Ref. 4.

by the same mechanism, reversing the domain-wall movement.

Since the net Ho–Ir interaction energy is $\pm h_{loc}\sum_i \sigma_i$, the energetic pressure between the two domains is $2h_{loc}\langle\sigma\rangle$ per iridium ion, which is plotted in Fig. 5.9. For small external fields, $\langle\sigma\rangle$ is controlled by the Ho–Ir coupling and, as such, it favours the existing domain locally. While type-A domains become energetically favourable in equilibrium for arbitrarily small positive fields, they can only grow if holmium moments rearrange inside a type-B domain to favour them locally: for small fields, this only happens by slow [272] thermal fluctuations of the moments, which slows down domain-wall movement substantially. Above a certain field (around 1.9 T at 1.8 K), however, $\langle\sigma\rangle$ becomes positive inside type-B domains, too: after this point, the growth of type-A domains is favourable everywhere, so it can occur via ultrafast iridium dynamics, only hindered by domain-wall pinning. This matches the rapid widening of hysteresis loops shown in Fig. 5.7(b) for $\mu_0 H_{max} \gtrsim 2$ T.

The experimental hysteresis loop closes around 2.9 T, well before the saturation of energetic pressure in Fig. 5.9. This suggests that the domain-wall pinning overcome by

the holmium-mediated energetic pressure is due to weak pinning sites (or potentially self-pinning of the complex magnetic structure), which dissipate net pinning energies on the order of a kelvin per iridium ion as the domain wall moves. This also helps explain the absence of long tails in the experimental hysteresis curves [cf. Fig. 5.6(c,d)]: the critical field increases only slightly with temperature [Fig. 5.9(b)], causing the experimental hysteresis curve to close abruptly at slightly higher fields.

Such a weak pinning, however, does not explain the fact that the ratio of iridium domains appears to saturate around 70 : 30 rather than 100 : 0. This requires additional rare but strong pinning sites that give rise to relaxation time scales longer than the experimental ones, which prevent the domain distribution from becoming any more polarised. I believe that these sites and the weak pinning discussed above are due to qualitatively different mechanisms, for instance impurities and self-pinning, respectively.

5.3.3 Resistivity signature of monopoles

Apart from field-polarised limits, the large holmium moments remain disordered in $Ho_2Ir_2O_7$: iridium conduction electrons may scatter off this disorder. I discuss two scattering processes, one involving the magnetic fields of the moments, the other due to magnetostriction effects [277]; I argue that the strength of both is proportional to the density of monopoles in the holmium spin ice, consistent with the similarity of the experimental resistivity and numerical monopole-density hystereses in Fig. 5.6(c,d). Finally, I demonstrate quantitatively that these scattering processes are able to account for the experimental resistivity hysteresis.

Magnetic scattering

The dominant interaction between the iridium conduction electrons and the localised holmium moments is the Zeeman coupling between the magnetic moments of the former and the magnetic field due the latter. Therefore, estimating the magnetic scattering rate requires accounting for the magnetic field patterns due to the disordered, yet strongly correlated holmium spins.

This task is greatly simplified by the dumbbell model of spin ice [62]: the dipole moment of each holmium ion is represented by two oppositely charged magnetic monopoles in the centres of the two tetrahedra the spin belongs to (§1.3.2). This way, most of the magnetic field inside the sample can be accounted for in terms of the Coulomb field of spin-ice monopoles living inside the tetrahedra that violate the 2120 rule [62, 79, 80]. The true field profile differs from this approximation by the field of the quadrupolar and higher-order moments that arise as the difference between the dumbbells and the phys-

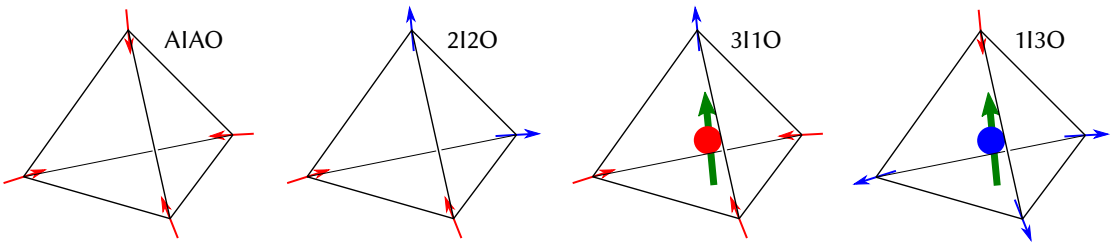

Figure 5.10. Magnetoelectric charge redistribution on holmium moments in spin ice and $Ho_2Ir_2O_7$ [277]. On AIAO or 2I2O tetrahedra, the charge redistribution produces no net electric dipole moment; on single monopoles (3I1O or 1I3O tetrahedra), an electric dipole develops pointing towards the minority spin.

ical holmium moments: these fields, however, decay at least as $1/R^4$ compared to the $1/R^2$ Coulomb field, so their effect is negligible even on the closest iridium sites.

Therefore, to a good approximation, the effect of spin-ice magnetism on the conduction electrons is described by the effective dipole–charge potential

$$V_{\mathrm{m}}(\vec{r}) = -\frac{\mu_0 q}{4\pi}\frac{\vec{\mu}\cdot\vec{r}}{r^3}, \tag{5.3}$$

where $q = 2\mu_{\mathrm{Ho}}/a_{\mathrm{d}}$ is the charge of a single magnetic monopole and the magnetic moment $\vec{\mu}$ of conduction electrons is on the order of μ_{B}. The interaction strength is captured by the force constant $\mu_0\mu_{\mathrm{B}}q/(4\pi) \approx 0.24$ meVÅ^2.

Electric scattering

Spin-ice materials like $Dy_2Ti_2O_7$ are multiferroic, that is, externally applied magnetic fields cause electric polarisation and vice versa [274–276]. A mechanism for these effects was proposed by Khomskii [277], who used a Hubbard-type model to describe the holmium *f*-electrons and found that the formation of local magnetic moments leads to a spontaneous redistribution of electrons between different atoms. The resulting charge redistribution may then introduce a net electric polarisation in the sample. Ref. 277 found that 2I2O, all-in, or all-out tetrahedra develop no net dipole moment; single monopoles, on the other hand, gain an electric dipole that points towards the minority spin of the tetrahedron (Fig. 5.10). Conduction electrons will then scatter off these dipoles through the *electric* charge–dipole potential

$$V_{\mathrm{e}}(\vec{r}) = -\frac{e}{4\pi\varepsilon_0}\frac{\vec{p}\cdot\vec{r}}{r^3}. \tag{5.4}$$

The magnitude of the emergent electric dipoles p is not known exactly; based on the polarisation measurements of Ref. 276, and magnetostriction effects in general, $p \sim 10^{-4}\, e\text{Å}$ appears plausible. Using this figure, $ep/(4\pi\varepsilon_0) \approx 1.4$ meVÅ^2, within one order of magnitude to the magnetic force constant derived above. This shows that both effects are similarly important in accounting for scattering off spin-ice monopoles.

Effect of monopoles on resistivity

In order to estimate the contribution of scattering off monopoles to the resistivity, we need to make some working assumptions about the nature of the conduction electrons, which is not very well understood. The experimentally observed very slow increase of resistivity below the metal–insulator transition suggests that the system is not a band insulator (especially compared to $Nd_2Ir_2O_7$ [258, 279]), but rather a (semi)metal or a heavily doped semiconductor [262], or perhaps not even a Fermi liquid [266]. For simplicity, however, we assume the following:

- The conduction electrons form a (small) metallic Fermi surface at the bottom of a quadratic conduction band: $E(\vec{k}) = \hbar^2 k^2/2m$, where the effective mass m is close to the bare electron mass.

- Monopoles are dilute enough that electrons only scatter off one at a time.

- The orientation of conduction-electron spins and electric dipole moments can be taken as independently and uniformly distributed on the unit sphere. This likely holds for magnetic scattering, as the ordered iridium moments are small [267] and their onset coincides with the metal–insulator transition, both of which suggests that the remaining conduction electrons are not spin-polarised. For electric scattering, the dipoles are oriented towards the minority spins of 3in and 1in tetrahedra [277]. While these are constrained to the eight ⟨111⟩ directions, numerical evidence suggests that they are essentially uncorrelated, apart from the [111] field-polarised limit.

- Electric and magnetic scattering are uncorrelated: this is plausible as the relevant dipole moments are due to the holmium and iridium atoms, respectively.

Several of these are rather crude approximations; as such, we shall only look for an order of magnitude estimate of the resistivity.

Let us now consider scattering due to a generic dipole–charge potential. The transi-

tion matrix elements are given by its Fourier transform:

$$V(\vec{r}) = -C\,\frac{\hat{p}\cdot\vec{r}}{r^3} \implies V(\vec{q}) = \int \mathrm{d}^3 r\, e^{-i\vec{q}\cdot\vec{r}}\, V(\vec{r}) = 4\pi i C\,\frac{\hat{p}\cdot\vec{q}}{q^2}, \tag{5.5}$$

where $\hat{p}$ is the unit vector parallel to the (electric or magnetic) dipole and C is a generic coupling constant, equal to $ep/(4\pi\varepsilon_0)$ and $\mu_0\mu_\mathrm{B}q/(4\pi)$ in the electric and magnetic cases, respectively. Now, the scattering rate follows from general scattering theory [280] as

$$\tau^{-1} = \int \frac{\mathrm{d}^3 k'}{(2\pi)^3} W_{k,k'}(1 - \hat{k}\cdot\hat{k}'); \tag{5.6}$$

$$W_{k,k'} = \frac{2\pi}{\hbar} n_\mathrm{mp}\,\delta\left(E(\vec{k}) - E(\vec{k}')\right)\left|V(\vec{k}' - \vec{k})\right|^2, \tag{5.7}$$

where n_mp is the number density of magnetic monopoles. Given the assumptions above, the integral (5.6) can be evaluated to

$$\tau^{-1}(k) = \frac{8\pi m^2 C^2}{3\hbar^4 k^2} n_\mathrm{mp} v, \tag{5.8}$$

where the group velocity is $v = \mathrm{d}E/(\hbar\mathrm{d}k) = \hbar k/m$. This shows that the force constants C for the two scattering channels are to be added in quadrature; for simplicity, we take a single $C \sim 1$ meVÅ^2. The contribution to resistivity can now be estimated from the Drude model:

$$\Delta\rho = \frac{m}{n_\mathrm{e} e^2}\langle\tau^{-1}\rangle = \frac{4\pi m^2 C^2 n_\mathrm{mp}}{\hbar^3 k_\mathrm{F} n_\mathrm{e} e^2}, \tag{5.9}$$

where n_e is the number density of conduction electrons and $k_\mathrm{F} = (3\pi^2 n_\mathrm{e})^{1/3}$ is the Fermi wave vector, still assuming a quadratic dispersion.

Discussion

Equation (5.9) shows that the contributions of magnetic and magnetoelectric scattering are proportional to monopole density. This is consistent with the very similar shapes[†] of the measured resistance-hysteresis curves in Fig. 5.6(c) and the hysteresis in monopole density arising as a result of iridium-domain redistribution by the [111] field: The latter is plotted in Fig. 5.6(d), using the same 70 : 30 and 30 : 70 splits between

[†]There are three main discrepancies between the two curves: (i) The experimental curves show an additional negative magnetoresistance at small fields, which is likely due to the partial ordering of the residual magnetic quadrupole moments and the corresponding reduction in magnetic scattering. (ii) At high fields, the monopole density rises, unlike the resistivity: this is due to the full ordering of the holmium moments that reduces scattering below what is expected based on (5.9). (iii) At higher temperatures, the resistivity hysteresis loops close at substantially smaller fields: this is because the domain ratio is swept quickly from 30 : 70 to 70 : 30 once the energetic pressure makes this favourable (§5.3.2).

type-A and B domains for the up- and downsweeps as Fig. 5.6(b) (§5.3.1).

Furthermore, the resistivity contribution (5.9) can realistically account for the full width of the magnetoresistance hysteresis, which is about $\delta R \approx 5\ \mathrm{m}\Omega$ at its widest point at $T = 2\ \mathrm{K}$. Given the sample size $\ell \approx 200\ \mu\mathrm{m}$, this corresponds to a resistivity hysteresis $\delta\rho \approx 10^{-6}\ \Omega\,\mathrm{m}$. The corresponding width of the monopole-density hysteresis in our simulations is $\delta n_{\mathrm{mp}} \approx 5\ \mathrm{nm}^{-3}$. Together with $C \approx 1\ \mathrm{meV\AA}^2$, we can invert (5.9) to obtain the carrier-density estimate $n_{\mathrm{e}} \sim 10^{17}/\mathrm{cm}^3$. This is some six orders of magnitude below the carrier density of elemental metals, a sensible figure for a badly conducting semimetal like $Ho_2Ir_2O_7$. Indeed, the corresponding Fermi energy is on the order of 10 K, potentially consistent with a thermally populated, insulating conduction band.

5.4 $Ho_2Ir_2O_7$ in a [110] magnetic field

Standing in the middle of the road is very dangerous;
you get knocked down by the traffic from both sides.
— Margaret Thatcher

Both the ordered 2I2O state and the 3I1O–1I3O monopole crystals found, respectively, in large [100] and [111] external fields are stabilised by a wide range of field directions [marked white and red/blue, respectively, in Fig. 5.11(b)] around the precise ⟨100⟩ and ⟨111⟩ directions. The ⟨110⟩ crystallographic axes sit in the corners of these ranges: as two of the ⟨111⟩ local Ising axes are perpendicular to them, [110] fields cannot polarise two spins of each tetrahedron, so 2I2O, 3I1O, and 1I3O configurations are all allowed at arbitrarily high fields. Therefore, measurements in a [110] field are highly susceptible to the slightest misalignment thereof, which allows us to confirm several subtle consequences of the theoretical interpretation given in §5.3.

Magnetisation and resistivity measurements in approximately [110] oriented fields are shown in Fig. 5.11(a,c). No hysteresis can be noticed in the magnetisation; in sharp contrast, magnetoresistance measurements show substantial hysteresis loops with a very similar shape to those in [111] fields [Fig. 5.6(c)]. Interestingly, these loops open *after* the magnetisation curve has almost fully saturated.

To interpret these findings, it is crucial to remember that in a real experiment, the external field never points exactly along the [110] axis, so its direction falls into one of the ordered domains in Fig. 5.11(b). However, as long as the field direction is close to [110], saturation towards the fully polarised state takes place in two stages:

- First, the two spins not perpendicular to the [110] direction are polarised by the dominant field component. In the measured magnetisation, this leads to essentially

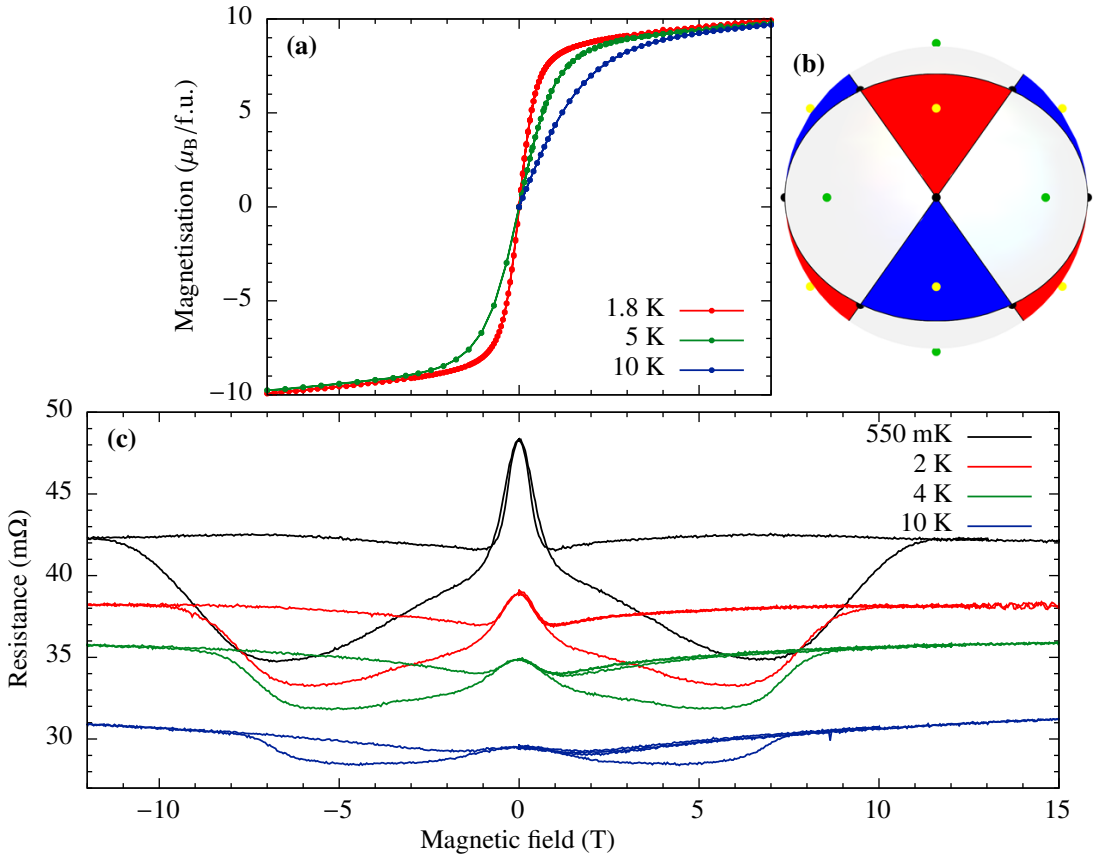

Figure 5.11. Magnetisation **(a)** and resistance **(c)** measurements on $Ho_2Ir_2O_7$ in an external [110] field. The magnetisation saturates somewhat above the expected 4.1 μ_B/Ho [270, 271] at fields around 1 T, without opening any visible hysteresis at any point. Large resistance-hysteresis loops of similar shape to those in Fig. 5.6(c) open above this field and remain open up to about 10 T.
(b) High-field polarised state of the holmium moments as a function of field direction. For almost all orientations, this is either an ordered 2120 state (white) consistent with a ⟨100⟩ (green dots) field, or a 3110–1130 monopole crystal (red and blue) brought about by a large ⟨111⟩ (yellow dots) field. ⟨110⟩ (black dots) fields do not couple to two spins on each tetrahedron, allowing for states consistent with any of these orders.
Figure based on Ref. 4.

full saturation at magnetisation $M^{\text{sat}}_{[110]} \approx 4.1\ \mu_B$/Ho [270, 271].

- At substantially higher fields (depending on the accuracy of [110] alignment), the two spins unconstrained by the [110] component get polarised, too. In this field range, the physics is determined by that in [100] or [111] fields, depending on which domain in Fig. 5.11(b) the field direction falls into.

In particular, the resistivity hysteresis shown in Fig. 5.11(c) is consistent with a [110] field canted slightly towards the [111] direction [i.e., pointing into a red/blue domain

in Fig. 5.12(b)]. It does not open up below 1–2 T, by which point the magnetisation has almost saturated. Afterwards, as all holmium moments saturate into the monopole crystal consistent with a [111] field, they exert the same energetic pressure on iridium AIAO domains that was discussed in §5.3.2, leading to a hysteretic redistribution of the latter. As a result, the magnetoresistance (which mostly depends on monopole density) shows a similar hysteresis as in a [111] field, but shifted to higher fields (2–10 T). On the other hand, the steady but non-hysteretic increase of magnetisation in Fig. 5.12(a) after "saturation" around 1 T is consistent with a [110] field canted towards the [100] direction (i.e., pointing into a white domain). This is possible because the two measurements were performed on two different samples, mounted independently.

We also considered theoretically the effect of rotating the field direction continuously through a ⟨110⟩ axis in order to demonstrate the robustness of the plastic domain-deformation mechanism proposed in §5.3.2. Specifically, in a field canted significantly towards a [111] direction [i.e., well inside a red domain in Fig. 5.11(b)], the hysteretic behaviour is not very different from that at ideal [111] alignment: in particular, a large hysteresis loop in resistance measurements is expected to open away from zero field. Upon approaching the [110] direction, the hysteresis would move to higher fields as the force acting on two spins in each tetrahedron is reduced. Eventually, the field necessary to fully open the hysteresis becomes larger than what is accessible in the experiment: at this point, the hysteresis loop becomes narrower [cf. Fig. 5.7(b)]. Afterwards, the field direction gets so close to the [110] axis that the hysteresis can no longer open even in a 15 T field. Further rotation carries the field direction from a red to a blue region, allowing the hysteresis to reopen once the field develops a sufficiently large $[11\bar{1}]$ component.

These expectations are all borne out by the magnetoresistance measurements shown in Fig. 5.12, which have been carried out specifically to check the theoretical predictions above. Beyond that discussion, an interesting feature captured in the top panel is a marked asymmetry of the magnetoresistance curve between positive and negative fields in the region (spanning about 5° of rotation angle) without any observable hysteresis. Initially, the resistivity is reduced more in positive fields, but after further rotation, the asymmetry switches sides. This is because the iridium AIAO domains remain unequally distributed after the last field sweep that did open a hysteresis, which pushed them to favour negative [111] fields. Therefore, the monopole density gets depleted in positive fields, causing a larger reduction of resistivity, even if the energetic pressure is too small to open any hysteresis. As the field direction is rotated through the [110] axis [from a red to a blue region in Fig. 5.11(b)], the direction of this asymmetry is reversed, since positive $[11\bar{1}]$ fields favour the same monopole crystal as negative [111] ones.

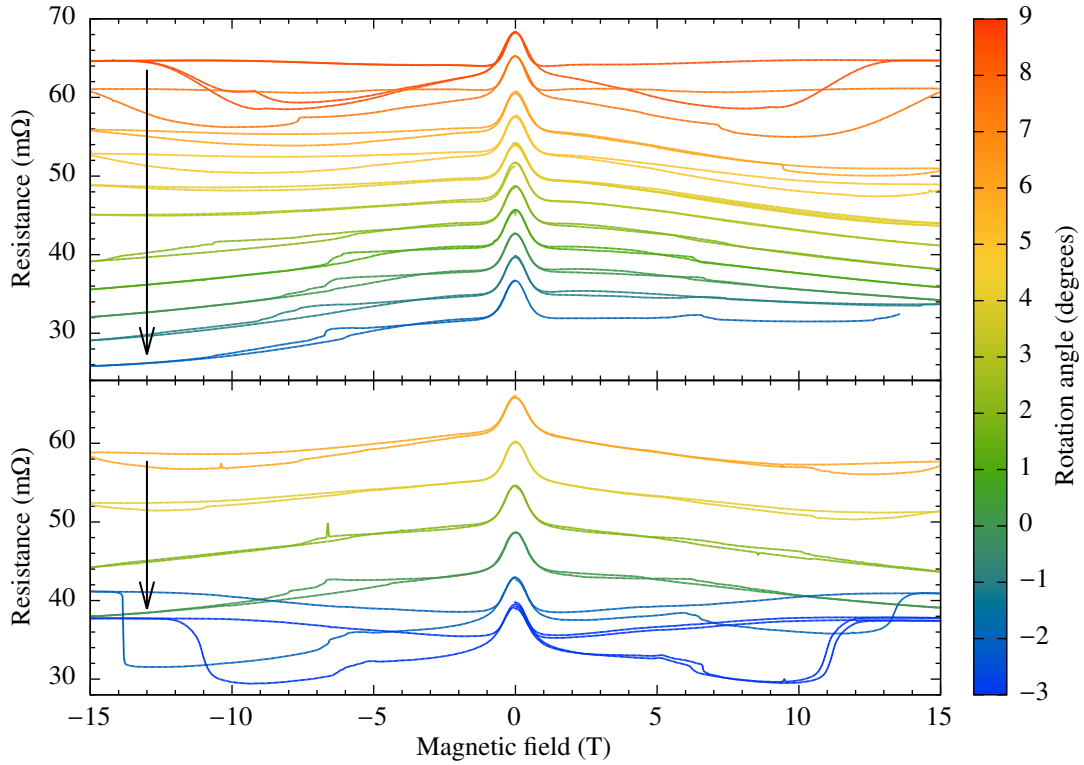

Figure 5.12. Magnetoresistance measurements on a sample rotated in an external field near the [110] direction. The rotation angle is given with respect to an arbitrary reference point along the route between (approximately) the [111] and [11$\bar{1}$] axes, passing near the [110] direction; the sample was rotated from higher to lower angles (black arrow). At high angles, a hysteresis loop similar to that of Fig. 5.6(c) opens. It remains closed at lower angles, but the magnetoresistance develops significant asymmetry: initially, resistance reduces more strongly in positive fields, but at negative angles, it becomes more prominent for negative fields. At the lowest negative angles, the hysteresis opens again. For clarity, the curves are shifted by 4 mΩ/degree. A contact was broken and repaired between the measurements in the two panels, so resistances and rotation angles cannot be compared directly. Figure based on Ref. 4.

5.5 Conclusion

For scientific endeavor is a natural whole the parts of which mutually support one another in a way which, to be sure, no one can anticipate.

— Albert Einstein

Since the earliest indirect evidence of magnetic monopoles in spin ice [62], much effort has been devoted to their direct detection and characterisation and, in particular, to measuring their density in experiments. The monopole density relates to both thermodynamic and dynamical properties of these materials; indeed, several quantities have been

proposed as measurement proxy, for example specific heat [65], spin correlations [65–67], and magnetic susceptibility and noise [73, 80, 81, 272, 281–283]. Our results show that magnetoresistance in the spin-ice compound $Ho_2Ir_2O_7$ is strongly linked to the concentration of magnetic monopoles in a way that holds promise to develop a readily measurable and versatile experimental indicator of their density. While we investigated the case of iridates, it will be interesting to see whether other spin-ice compounds can be found with small enough gap that a similar approach can be used. Since the relation between resistance and monopole density does not hinge on the magnetism of iridium ions, it could be used in principle to measure the monopole density in other spin-ice systems where only rare-earth magnetism is present, such as $Dy_2Ti_2O_7$ and $Ho_2Ti_2O_7$.

Furthermore, our results demonstrate that the magnetisation and magnetoresistance of $Ho_2Ir_2O_7$ becomes hysteretic in an applied [111] magnetic field due to plastic deformation of the antiferromagnetic iridium domain walls. Antiferromagnetic domains are promising building blocks for future spintronic devices as they do not produce stray magnetic fields, are not susceptible to the same, and have ultrafast spin dynamics [252]. However, the manipulation of antiferromagnetic domain walls is challenging precisely due to their zero net magnetisation, which requires a staggered field to interact with the alternating moments [252]. This is circumvented in $Ho_2Ir_2O_7$ by the interplay between the iridium domains and the frustrated ferromagnetism of the large holmium moments, which acts to drive the motion of the domain walls. This provides highly reproducible, robust, and precise control over the antiferromagnetic domains using an external magnetic field; importantly, the domains remain stable against weak magnetic fields, as the coupling mechanism becomes relevant only when the holmium spins become nearly polarised. While driving the domain walls in $Ho_2Ir_2O_7$ is only possible at low temperatures and relatively large fields, it serves as an example of a wider class of materials in which magnetic-field control of antiferromagnetic domains is possible: these contain two, strongly coupled magnetic species, one with large, frustrated moments, the other realising robust long-range antiferromagnetic order.

We believe that the plastic deformation of iridium domains are controlled by a distribution of pinning energies for the domain walls. While such plastic behaviour is generally expected for magnetic domain walls [273], understanding and modelling it in these materials is an open and interesting question of direct relevance to potential applications, such as those in spintronics.

Finally, our results highlight the inherent interplay between the antiferromagnetic iridium order and frustrated rare-earth ferromagnetism present in spin-ice iridates, which generates a rich and exciting playground for the study of complex and out-of-equilibrium behaviour extending beyond that offered by the hitherto better-known titanate spin ices.

6

Localisation in one- and two-dimensional quasicrystals

In all chaos there is a cosmos, in all disorder a secret order.

— Carl Gustav Jung

It was long believed that all long-range ordered structures are crystalline, that is, they consist of a periodic repetition of some unit cell. The first counterexamples to this notion were found among tilings of the plane [284], the best known of which are the *Penrose* [285] and *Ammann–Beenker* tilings [286]. Interest in quasiperiodicity within the physics community was sparked by the discovery of *quasicrystals* by Shechtman [287] and the equivalence between Landau levels on two-dimensional lattices and a one-dimensional quasiperiodic chains [288–290]. Recently, quasiperiodic structures became popular in ultracold-atom experiments as a proxy for random potentials in the study of disordered quantum gases, Bose glasses, and many-body localisation (MBL), as they can conveniently be realised by superimposing two incommensurate optical lattices [291–300]. Quasiperiodic tilings also lie at the heart of recent results in the study of quantum complexity, such as the proof of the undecidability of the spectral gap [301].

The single-particle behaviour of quasicrystals and other quasiperiodic systems exhibits a range of unusual features that are also important to understand their many-body physics. It is well known that arbitrarily weak uncorrelated disorder causes Anderson localisation in one dimension [302]; by contrast, the fate of eigenstates in a quasiperiodic system depends strongly on the strength and functional form of the disorder. For instance, *Fibonacci word models* are tight-binding Hamiltonians of the form

$$H_t = -\sum_j t_j \left(a_j^\dagger a_{j+1} + \text{H.c.}\right); \tag{6.1a}$$

$$H_V = -t \sum_j \left(a_j^\dagger a_{j+1} + \text{H.c.}\right) + \sum_j V_j a_j^\dagger a_j, \tag{6.1b}$$

where the hopping matrix elements t_j and on-site potentials V_j alternate between two distinct values according to a quasiperiodic sequence called the Fibonacci word [303–305]. Remarkably, for a wide range of parameters, the eigenstates of (6.1) have a multifractal structure, usually only seen at critical points [303–307]. Quasiperiodic models also exhibit Anderson localisation transitions in one dimension [308–310] and other unusual transport properties in higher dimensions [311, 312]. Furthermore, quasiperiodic potentials contain no rare regions in the usual sense, that is, patches in which the local disorder is by chance substantially lower than on average. Such regions give rise to Griffiths effects [313], which are expected to substantially affect MBL in disordered one-dimensional systems [314, 315] and destabilise MBL completely in higher dimensions [316]. Quasiperiodicity might therefore prove essential to stabilising MBL in higher dimensions.

Quasiperiodic systems also inherit fascinating topological properties from higher-dimensional periodic parent Hamiltonians [317] from which they can be derived using cut-and-project methods [318]. As an example, two-dimensional quasicrystals were observed to exhibit the four-dimensional integer quantum Hall effect [319, 320].

In §6.1, I discuss renormalisation-group (RG) studies [5] of the Anderson localisation transition in the *Aubry–André (AA) model* [308], a one-dimensional quasiperiodically modulated tight-binding chain. RG methods have been used extensively to study critical dynamics in quasicrystals [306, 307, 321–328]. In contrast with common RG methods (e.g., the momentum-shell approach [139]), however, the structure introduced by quasiperiodicity does not allow renormalisation by an arbitrary or infinitesimal amount. As a result, quasicrystals are only amenable to *discrete RG methods*, where the renormalisation of length scales is matched to the period of quasiperiodic modulation [306, 307, 321, 324–326, 329–333]. I found that this discrete RG flow prevents the formation of a scaling regime described by universal power laws; indeed, critical exponents near the second-order localisation transition cannot be defined for almost all values of the irrational number β that describes the quasiperiodicity of the model. Nevertheless, I found that the quantum localisation transitions of a number of models sharing the same value of β display identical behaviour in the limit of large correlation lengths: this is the hallmark of a novel type of *non-power-law universality*, only captured by the detailed dependence of observables on length scales like the correlation length ξ.

In §6.2, I introduce a two-dimensional quasiperiodic Hamiltonian on the square lattice [6], which admits a duality transformation similar to that of the 1D Aubry–André model. In the latter case, this duality causes a localisation transition in all eigenstates at a critical strength of the quasiperiodic potential; by contrast, the new *2DAA model* exhibits a mixed spectrum of interspersed localised and partially extended eigenstates for arbitrarily strong quasiperiodic modulation. We found that the partially extended states are

confined to a set of special lattice lines with unusually weak modulation, while localised eigenstates live in the rest of the lattice. This geometric separation prevents hybridisation between these states and thus the formation of clean mobility edges [334]. These "rare lines" provide a fascinating alternative to both the two-dimensional rare regions in typical disordered systems and the complete absence of rare regions normally envisaged in the quasiperiodic case, with a potential for novel MBL properties.

6.1 Non-power-law criticality in the Aubry–André model

More is different.

— Philip W. Anderson

The Aubry–André model [308] (also known as the *Harper model* [288–290]) is given by the quasiperiodically modulated one-dimensional tight-binding Hamiltonian

$$H = -J \sum_j \left(a_j^\dagger a_{j+1} + \text{H.c.} \right) - J\lambda \sum_j \cos(2\pi\beta j) a_j^\dagger a_j, \tag{6.2}$$

where $a_j^\dagger$ is a creation operator on site j, J is the hopping matrix element, and $\beta \notin \mathbb{Q}$ and λ are the incommensurate wave number and dimensionless amplitude of the modulation, respectively. Unlike 1D systems with uncorrelated disorder [302], this model displays a localisation transition at $\lambda = 2$: below this amplitude, all eigenstates are extended, above it, all are localised [308, 321–323, 335]. This can be understood considering *Aubry duality*: applying the Fourier transform

$$b_k = \frac{1}{\sqrt{N}} \sum_n \exp\left(2\pi i \beta k n\right) a_n \tag{6.3}$$

to (6.2), it turns into a momentum-space Aubry–André Hamiltonian with λ changed to $4/\lambda$, and all energies rescaled by a factor of $\lambda/2$ [308]; the transition point $\lambda = 2$ is the fixed point of this transformation with critical states that remain unchanged by (6.3).

6.1.1 Renormalisation group for 1D quasicrystals

The spectra of one-dimensional quasiperiodic systems is a hierarchical Cantor set [303, 305–307, 321, 322, 324, 340–342], that is, it consists of a succession of smaller and smaller gaps. The positions of these gaps are fixed by the gap-labelling theorem [343], meaning that the hierarchy of gaps is topologically protected between different quasiperiodic models described by the same irrational wave number β [344, 345]. In the case of tight-binding models, the spectrum is bounded and its entire structure is governed by the

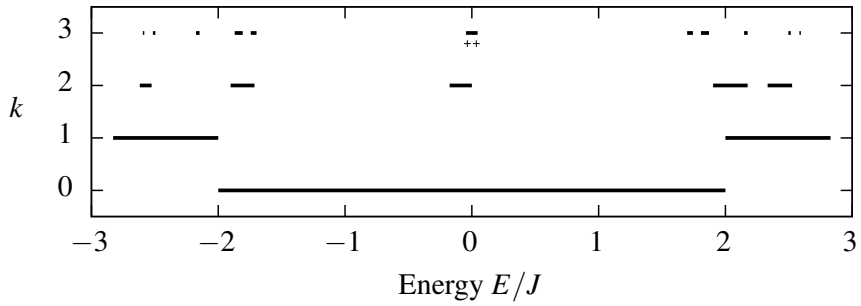

Figure 6.1. Spectrum of the Aubry–André model for different rational $M_k/N_k \approx [0;\overline{2}]$ at the critical point $\lambda = 2$. In each case, the spectrum consists of N_k bands, most of which are accounted for by splitting the N_{k-1} bands of the previous rational approximation into n_k narrower ones. Some additional bands appear due to the slight changes to the approximation of β [290, 336–339]. Double crosses denote a pair of bands with a very small gap, which is not resolved well in the plot. Figure taken from Ref. 5.

continued fraction expansion of β [305, 321, 346],

$$\beta = \frac{1}{n_1 + \beta_1} = \frac{1}{n_1 + \dfrac{1}{n_2 + \beta_2}} = \cdots = \frac{1}{n_1 + \dfrac{1}{n_2 + \dfrac{1}{\ddots}}} \equiv [0; n_1, n_2, n_3, \dots], \tag{6.4}$$

where the n_k are integers and the irrational residuals β_k are between 0 and 1. To demonstrate this, let us consider a β for which all n_k are very large; however, qualitative conclusions remain the same for all $n_k \geq 2$[†] [290, 338, 339]. Now, the structure of the spectrum can be described in terms of a sequence of periodic superlattices described by $M_k/N_k = [0; n_1, \dots, n_k]$, which are the closest rational approximations of β in the sense that [346]

$$|N_k\beta - M_k| < |N\beta - M| \qquad \forall M, N \in \mathbb{Z}, \quad 0 < N < N_k.$$

At the first step of this protocol, $M_1/N_1 = 1/n_1$: Bloch's theorem applies to the superlattice of period n_1, resulting in a spectrum with n_1 subbands of continuous dispersion. At the next step, the period of the superlattice and thus the number of subbands is $N_2 = n_1n_2 + 1 \approx n_1n_2$.[‡] Since the approximation to β changes very little, the spectrum

[†]The case of $n_k = 1$ is special. It implies $\beta_{k-1} > 1/2$, which can be replaced with $1 - \beta_{k-1}$ without changing the resulting structure: the first continued fraction term of this number is no longer 1.

[‡]Not all basis states of the system are accounted for in this approximation. This is compensated for by additional kth-order subbands that appear in the middle of each subband of order $(k-2)$ (Fig. 6.1) and are also described by the incommensurate ratio β_k [290, 338, 339]. The approximation, however, provides a good description of states near the edges of the spectrum [321].

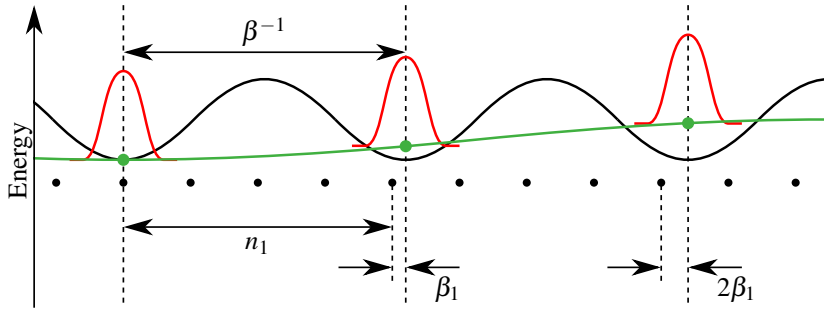

Figure 6.2. Cartoon of the renormalisation transformation of the lowest subband of the Aubry–André model. Black dots mark the sites of the original tight-binding chain: their quasiperiodic on-site potentials are indicated by the black line. In the first step of the transformation, every band gives rise to one effective Wannier state (red) for each period of the on-site potential. For each band, an effective tight-binding Hamiltonian can be written down using these Wannier states as renormalised lattice sites. Due to the incommensurability of the potential and the original lattice, the jth such Wannier state is shifted by $j\beta_1$ relative to the original lattice sites. As a result, the effective Hamiltonian is modulated quasiperiodically, with incommensurate ratio β_1 (green line). The procedure can then be repeated using the first-order Wannier states as lattice sites (green dots), introducing β_2, and so on indefinitely. Figure taken from Ref. 5.

is still dominated by the n_1 first-order bands, each of which splits into n_2 narrower subbands (Fig. 6.1). As further continued-fraction terms are taken into account, finer and finer subbands are formed by splitting existing subbands into n_k new ones at the kth step.

This hierarchical structure can be understood in terms of a discrete renormalisation-group procedure [5, 306, 307, 321, 322, 324]. Creating first-order subbands can be taken as renormalising length scales by a factor of $\beta^{-1} \approx n_1$: the new "lattice sites" correspond to approximate Wannier states located at each minimum of the quasiperiodic modulation (Fig. 6.2). Since the modulation period is incommensurate to the lattice spacing, the jth renormalised lattice site will have a phase shift $2\pi\beta_1 j$ compared to the original lattice sites. This results in a quasiperiodic modulation with incommensurate ratio β_1 in the effective Hamiltonian of each subband. These Hamiltonians can now be renormalised by a factor of $\beta_1^{-1} \approx n_2$, giving rise to second-order subbands modulated with the new incommensurate ratio β_2: repeating such steps indefinitely constructs the entire spectrum. In a critical Hamiltonian, the ground-state wave function can be described in terms of these Wannier states, resembling wave packets, generated at each step of the RG protocol; for β with periodic continued-fraction expansions, this leads to self-similar ground states (Fig. 6.3).

It is important to note that for generic β, the sequence of the n_k is arbitrary, so each

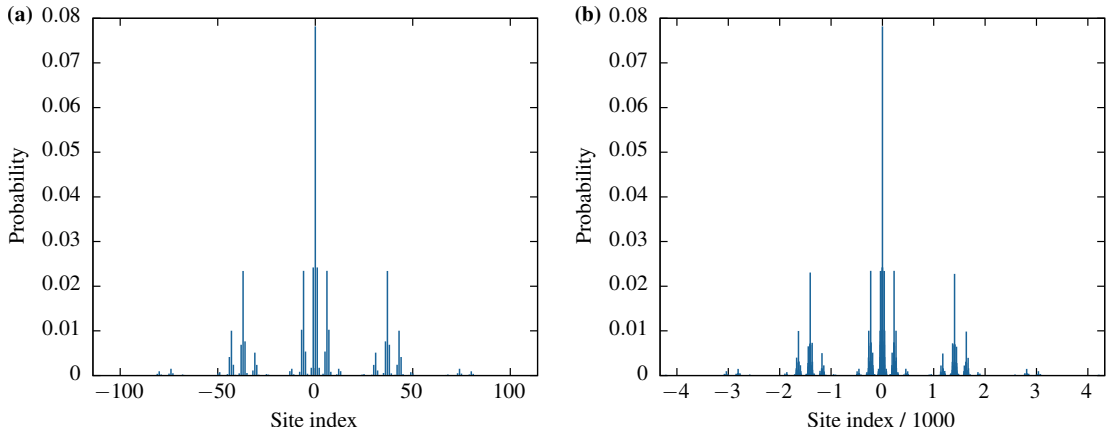

Figure 6.3. Ground-state probability distribution $|\psi_0(n)|^2$ in the Aubry–André model for $\lambda = 2$ and $\beta = 1405/8658 \approx [0; \overline{6}]$.
(a) The wave function can be constructed out of wave packets centred on each minimum of the quasiperiodic potential (period: $\beta^{-1} \approx 6.162$): these wave packets themselves form "higher-order wave packets" of size $\beta^{-2} \approx 37$, etc.
(b) At criticality ($\lambda = 2$), this construction can be continued indefinitely: the structure of the wave function is self-similar between length scales.

RG step is distinct from any other, leading to spectra and wave functions without self-similar features. However, if β is the root of a quadratic equation, its continued-fraction expansion is periodic [346]; now, all RG steps within one period can be combined into a single step which can be repeated indefinitely, building up features with (discrete) self-similarity. I am going to refer to such numbers as *quadratic irrationals,* and denote a periodic continued fraction expansion with an overbar (e.g. $[0; \overline{2, 3}] \equiv [0; 2, 3, 2, 3, \dots]$).

Details for the Aubry–André model

To explicitly construct the renormalised effective Hamiltonian of the Aubry–André model (6.2), we assume that all $n_k \gg 1$ and replace β with $1/n_1$ in the first approximation:

$$H = -J \sum_j \left(a_j^\dagger a_{j+1} + \text{H.c.} \right) - J\lambda \sum_j \cos\left[\frac{2\pi}{n_1} (j - \phi) \right] a_j^\dagger a_j, \tag{6.5}$$

where ϕ is a meaningful[†] spatial offset. Since this Hamiltonian is periodic, Bloch's theorem applies: the spectrum of (6.5) consists of n_1 narrow bands, each of which gives rise to a set of Wannier states separated by n_1 lattice sites. For large n_1, couplings beyond nearest neighbours can be ignored, giving the dispersion relation

$$E(\kappa, \phi) = E_0(\phi) - 2J'(\phi) \cos \kappa, \tag{6.6}$$

[†]If β were irrational, any such offset could be approximated arbitrarily well by $n\beta$ for some $n \in \mathbb{Z}$, so ϕ would only amount to shifting the origin.

where J' is the hopping between two neighbouring Wannier states, κ is the renormalised quasimomentum, and E_0 is the mean energy of the band. In principle, both E_0 and J' depend on the phase ϕ in (6.5). It turns out, however, that the modulation of E_0 is on the order of J' which is exponentially small [321]; likewise, the modulation of J' is on the order of second-neighbour hopping, which we neglect.

We now apply the Aubry duality transformation (6.3) to (6.5) [5, 321]. As β is now rational, Eq. (6.3) only generates n_1 distinct Fourier modes; the transformation can be made unitary by considering only n_1 real-space sites with twisted periodic boundary conditions:

$$H = -J\sum_{j=1}^{n_1}\left(a_j^\dagger a_{j-1}e^{i\kappa/n_1} + \text{H.c.}\right) - J\lambda\sum_{j=1}^{n_1}\cos\left[\frac{2\pi}{n_1}(j-\phi)\right]a_j^\dagger a_j \tag{6.7a}$$

$$= -\frac{J\lambda}{2}\left[\sum_{\ell=1}^{n_1}\left(b_\ell^\dagger b_{\ell-1}e^{2\pi i\phi/n_1} + \text{H.c.}\right) + \frac{4}{\lambda}\sum_{\ell=1}^{n_1}\cos\left\{\frac{2\pi}{n_1}\left(\ell - \frac{\kappa}{2\pi}\right)\right\}b_\ell^\dagger b_\ell\right], \tag{6.7b}$$

where $a_0 = a_{n_1}$ and $b_0 = b_{n_1}$. That is, the duality transformation exchanges the quasimomentum κ and the offset ϕ: combining this with (6.6) gives the full dispersion relation

$$E(\kappa,\phi) = E_0 - 2J'\cos\kappa - (J\lambda)'\cos(2\pi\phi); \tag{6.8}$$

$$(J\lambda)' = \frac{\lambda}{2}\times 2J'_{\text{dual}}. \tag{6.9}$$

In the original quasiperiodic Hamiltonian, however, β^{-1} differs from n_1 by a small irrational number β_1; therefore, the jth minimum of the potential is shifted by $j\beta_1$ from the original lattice sites (Fig. 6.2). Since $\beta_1 \ll 1$, Eq. (6.8) is still a good effective Hamiltonian, provided ϕ changes slowly as a function of position:

$$H = -2J'\cos\hat{p}' - (J\lambda)'\cos(2\pi\beta_1\hat{x}'). \tag{6.10}$$

Eq. (6.10) is an Aubry–André model of incommensurate ratio β_1 and renormalised potential and hopping terms. The same procedure can then be repeated with step sizes β_k^{-1} to construct the full spectrum.

Renormalisation of λ and J: the critical exponents ν and z

In the limit $n_k \gg 1$, the discreteness introduced by lattice sites can be ignored: treating wave packets as continuous, their wave functions and the hopping term J' can be estimated using a WKB approximation sufficiently far from $E = 0$ [5, 321, 347, 348]. These calculations show that the renormalisation of the dimensionless amplitude λ does not

depend on energy:

$$\lambda' = 2\left(\frac{\lambda}{2}\right)^{1/\beta}. \tag{6.11}$$

$\lambda = 2$ is an unstable fixed point of this RG transformation, that is, all $\lambda < 2$ renormalise to $\lambda = 0$ (extended phase) and all $\lambda > 2$ renormalise to $\lambda \to \infty$ (localised phase). Furthermore, the critical exponent of the reduced tuning parameter, $g = \log(\lambda/2)$ is $\nu = 1$ regardless of β: indeed, it can be shown [5, 308] that

$$\xi = \left|\log \frac{\lambda}{2}\right|^{-1} \tag{6.12}$$

for all energies and values of β.

The dynamics of a quantum phase transition is usually characterised using the *dynamical exponent* z defined through the lowest-lying energy gap ΔE of the system as

$$\Delta E \propto \xi^{-z} \implies z = -\lim_{\xi\to\infty} \frac{\log \Delta E(\xi)}{\log \xi}. \tag{6.13}$$

As the kth renormalisation step of the critical ($\lambda = 2$) Aubry–André model considers the rational approximation

$$\beta \approx [0; n_1, n_2, \ldots, n_k] \equiv \frac{M_k}{N_k}, \tag{6.14}$$

the correlation length is limited by the period of this approximation, N_k. Seeing that the renormalised J in this step is of the same order of magnitude as the lowest-lying band gap, z can be estimated through [5]

$$\frac{J^{(k)}}{J^{(k-1)}} \approx \frac{\Delta E(N_{k+1})}{\Delta E(N_k)} \equiv \left(\frac{N_{k+1}}{N_k}\right)^{-z_k} \simeq \beta_k^{z_k}. \tag{6.15}$$

That is, z_k only depends on the kth renormalisation step, which is fully governed by β_k. Remarkably, z_k does strongly depend on β_k: to leading order [5],

$$z_k \approx 1.166 \frac{\beta_k^{-1}}{\log(\beta_k^{-1})}. \tag{6.16}$$

This implies that the scaling of the energy gap ΔE changes between renormalisation steps, and so it never reaches a conventional scaling regime described by a simple power law. For quadratic irrationals, the sequence of β_k, and thus of renormalisation steps, is periodic: by averaging over one such period, it is possible to find an average z which satisfies the limit (6.13); however, log-periodic oscillations persist around it to arbitrarily large length scales [329–333]. By contrast, if β is non-quadratic, there is no periodicity

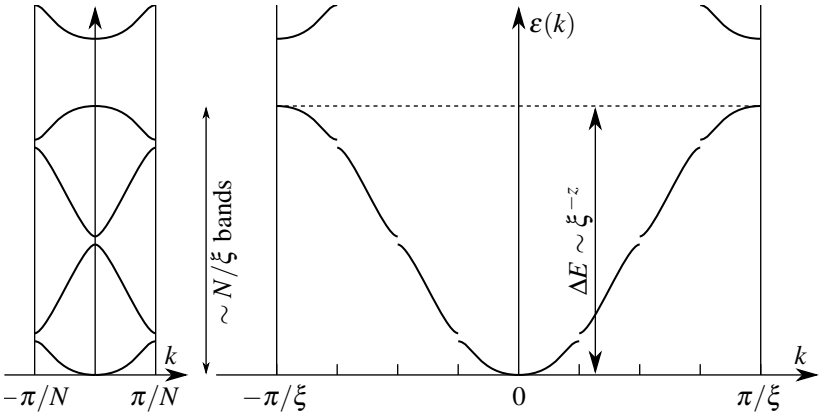

Figure 6.4. Cartoon of the extended lowest band of a periodic tight-binding model of period $N \gg \xi$. On length scales N_k above the correlation length, any effective potential is negligibly weak and results only in narrow avoided crossings. That is, the band structure is similar to that of an uninterrupted band of a model with period $\sim \xi$, as shown in the extended Brillouin zone on the right. The effective lowest band that governs quantum critical dynamics is thus $\sim 2\pi/\xi$ wide in momentum space and its width ΔE tends to a constant value as $N \gg \xi$: therefore, $\Gamma \sim \xi^2 \Delta E \sim \xi^{2-z}$. Figure taken from Ref. 5.

in the renormalisation steps, so the limit (6.13) generally does not exist.

6.1.2 Critical scaling near the localisation transition

We performed exact diagonalisation on the single-particle Hamiltonian (6.2). The behaviour of the incommensurate model was extrapolated from the sequence of rational approximations $M_k/N_k = [0; n_1, \ldots, n_k]$ to a given β, implemented with periodic boundary conditions. We evaluated the curvature of the lowest band,

$$\Gamma = \frac{1}{2Ja_0^2} \left. \frac{\mathrm{d}^2\varepsilon(k)}{\mathrm{d}k^2} \right|_{k=0} = \frac{1}{J} \frac{\hbar^2}{m_{\mathrm{eff}} a_0^2}, \tag{6.17}$$

where m_{eff} is the effective mass of particles near the bottom of the band and a_0 is the lattice spacing. In an extended phase, the single-particle dynamics becomes ballistic beyond a length scale, therefore, its effective mass tends to a finite value in the limit of an infinite system. Bands of a localised model, however, become completely flat, resulting in an infinite effective mass and thus zero curvature. That is, Γ for the rational approximations has a well-defined limit as $N_k \to \infty$, which can be used as an order parameter of the quantum localisation transition. In interacting many-particle systems, the appropriate generalisation of Γ gives the superfluid fraction or superfluid stiffness, which is widely used to analyse superfluid–insulator transitions [349–351].

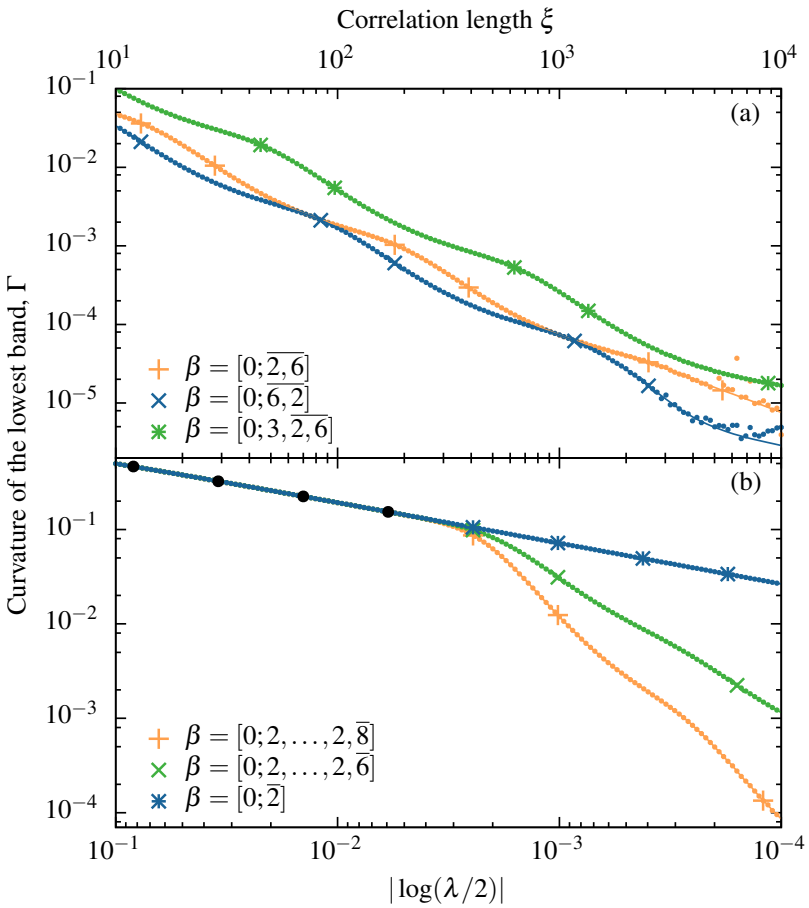

Figure 6.5. Curvature Γ of the lowest band as a function of the reduced tuning parameter $g = \log(\lambda/2) = \xi^{-1}$ [5] for several incommensurate ratios β. Small dots indicate all computed data points, large symbols appear at $g = 1/N_k$. Smoothing splines (thin solid lines) were added to two data sets to guide the eye. The period of the simulated superlattice for each curve is much greater than $\xi_{\max} = 10^4$.
(a) All β are quadratic irrationals: their continued-fraction expansions are identical except for the first few terms. The sequences $\{\beta_k\}$ that govern the fine structure of the spectrum are hence the same up to a shift, resulting in identical line shapes and effective critical exponents. The discrete nature of the underlying RG flow is manifest in log-periodic oscillations [329–333].
(b) The three values of β differ by less than a part in 10^6; however, their continued-fraction expansions diverge after the 8th term ($N_8 = 985$), resulting in identical scaling up to $|g| \gtrsim 1/985$ (black dots) followed by markedly different critical behaviour for $|g| \ll 1/985$.
Figure taken from Ref. 5.

The width ΔE of the lowest band is clearly related to its curvature, so Γ can be used to measure the dynamical exponent. In an extended phase, the *effective* lowest band has width $\sim 2\pi/\xi$ in reciprocal space because on larger length scales, the effective Hamiltonian is free electron-like with minor gaps introduced by a weak residual potential

(Fig. 6.4): as a result, the scaling of Γ is given by [5, 352]

$$\Gamma \sim \frac{\Delta E}{(2\pi/\xi)^2} \sim \xi^{2-z} \sim |\lambda - 2|^{\nu(z-2)}, \tag{6.18}$$

where the correlation length exponent ν is always 1 in the Aubry–André model, cf. (6.12).

To investigate the localisation transition of the Aubry–André model, Γ was calculated near $\lambda = 2$ for several different values of β; the results are plotted in Fig. 6.5. The rational approximations to β were chosen such that the periods of the superlattices simulated were much larger than the largest correlation length considered, $\xi_{\max} = 10^4$. As predicted, quadratic irrational values of β give rise to overall power-law critical behaviour with log-periodic oscillations [Fig. 6.5(a)]: the line shape of these oscillations is identical for different β with the same continued fraction period. This is readily explained by the scaling behaviour at $\xi \approx N_k$ being governed by the kth renormalisation step and hence by β_k only. This notion is further confirmed by the data in Fig. 6.5(b): at small correlation lengths, the RG flow is determined by the first few β_k that are all very close to one another, resulting in indistinguishable scaling curves; at higher ξ, diverging continued-fraction terms take over and lead to qualitatively different large-scale behaviour. This demonstrates how non-quadratic irrational β, whose continued-fraction terms do not follow a simple pattern, give rise to arbitrarily complex non-power-law critical behaviour.

Non-power-law universality

We studied localisation transitions in two quasiperiodic Hamiltonians beyond the Aubry–André model. The first is the generalised AA model

$$H = -\sum_n \left[J + J\mu\cos\left\{2\pi\beta\left(n+\tfrac{1}{2}\right)+\phi\right\}\right]\left(a_n^\dagger a_{n+1} + \text{H.c.}\right) - J\lambda\sum_n \cos(2\pi\beta n) a_n^\dagger a_n. \tag{6.19}$$

Similar to the Aubry–André model, eigenstates of (6.19) are either all extended or all localised. For $\phi \neq 0$, the boundary of extended and localised phases is given by [310]

$$\sum_\pm \sqrt{(\lambda/2)^2 \pm \lambda\mu\cos\phi + \mu^2} = 2. \tag{6.20}$$

For $\phi = 0$, the phase diagram consists of extended ($\lambda < 2; \mu < 1$), localised ($\lambda > 2, 2\mu$), and critical ($2\mu > 2, \lambda$) phases [309, 310]. As examples, localisation transitions along the

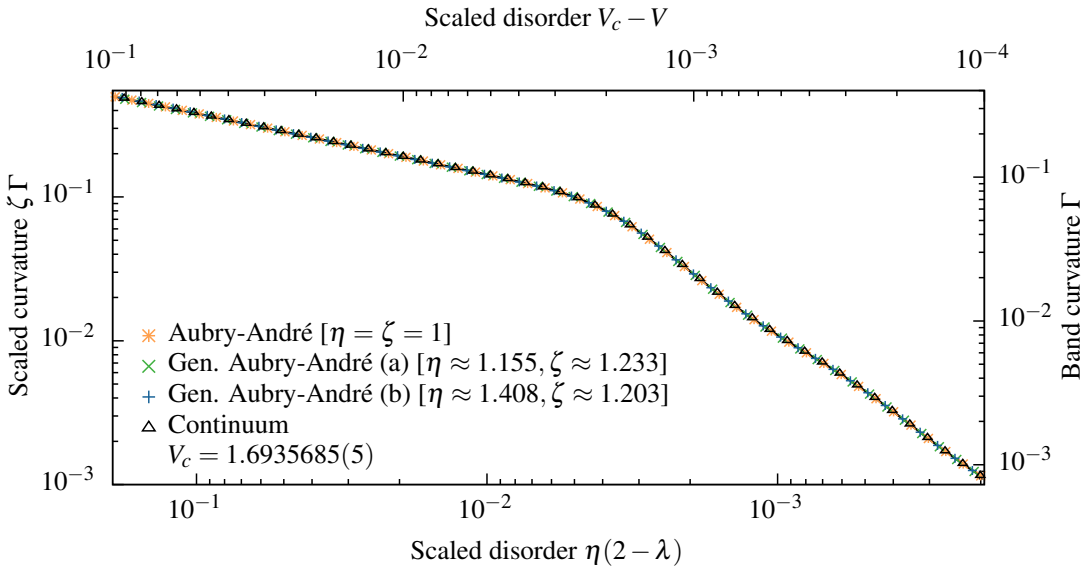

Figure 6.6. Curvature of the lowest band in the Aubry–André model, the generalised AA model (6.19, 6.21) (bottom and left axes), and the continuum model (6.22) with $V_1 = V_2$ (top and right axes) as a function of the reduced tuning parameters $\lambda - 2$ and $V - V_c$, respectively, for $\beta = [0; 2, \ldots, 2, \overline{6}]$. Apart from overall rescaling, the scaling behaviour of all models are equivalent, suggesting that they belong to the same universality class, which, however, is not properly described by power-law scaling. Figure based on Ref. 5.

following paths were considered:

$$\phi = 0 \qquad \mu = 1/2 \qquad \lambda \approx 2; \tag{6.21a}$$

$$\phi = \pi/2 \qquad 2\mu = \lambda = \Lambda/\sqrt{2} \qquad \Lambda \approx 2. \tag{6.21b}$$

The second model is the continuum quasiperiodic Hamiltonian

$$H = \frac{\hat{p}^2}{2m} + V_1 \cos^2(kx) + V_2 \cos^2(\beta kx), \tag{6.22}$$

which reproduces the Aubry–André model in the limit $V_1 \gg E_\mathrm{r} \gg V_2$, where the recoil energy $E_\mathrm{r} = \hbar^2 k^2/2m$ is the typical kinetic energy scale of the system. In addition to this limit, we studied the case of equal lattice depths $V_1 = V_2 = VE_\mathrm{r}/2$. Periodic approximations to the Hamiltonian (6.22) were implemented in momentum space and the curvature of the lowest band was calculated by exact diagonalisation using a formula adapted from (6.17) [349, 350]:

$$\Gamma = \frac{\pi^2}{E_\mathrm{r}} \frac{E_\Theta - E_0}{(\Theta/N)^2}\bigg|_{\Theta \to 0} = \frac{m}{m_\mathrm{eff}}. \tag{6.23}$$

A quantum localisation transition was observed for all values of the incommensurate ratio β we tested at a β-dependent critical V_c. Unlike the generalised Aubry–André model, however, its spectrum is unbounded, and several mobility edges appear in the spectrum of excited states. Nevertheless, we expect that the ground state has a hierarchical structure similar to that shown in Fig. 6.3.

Γ was evaluated for the Aubry–André model, the generalised AA model (6.19, 6.21) and the continuum model (6.22) near their ground-state transitions for $\beta = 16110/38893 \approx [0; 2, \ldots, 2, \overline{6}]$ and plotted in Fig. 6.6. Although there is no power-law scaling attained at large ξ, the curves collapse on top of one another: points mapped together correspond to equivalent correlation lengths. The existence of such a mapping is the hallmark of universality between the ground-state localisation transitions of these models. This universality is, however, markedly different from the conventional notion of universal critical exponents that arise from a large-ξ scaling regime with continuous RG flow.

6.1.3 Single-particle quench dynamics

In addition to the ground-state localisation transition, we investigated the quantum dynamics of the generalised Aubry–André models (6.2, 6.19) at the critical point. Importantly, *all* eigenstates of these Hamiltonians go through the localisation transition simultaneously [308–310]; therefore, we anticipate that this dynamics will be neither ballistic nor localised on any length scale.

Specifically, we considered the dynamics of a quantum state localised on a single site x_0 after a sudden quench onto the critical Aubry–André Hamiltonian. The expansion of such a wave packet can be characterised through the time evolution of its pth moment:

$$\mu_p = \left\langle |x - x_0|^p \right\rangle; \qquad \ell_p = \mu_p^{1/p}, \tag{6.24}$$

where the exponent p can be an arbitrary positive real number. In a conventional critical system, $\ell_p \sim t^{1/z}$ because t is a characteristic time scale corresponding to the length scale $\ell_p(t)$ [353]. In this context, $\sigma = 1/z$ is commonly referred to as the *anomalous diffusion exponent*.

The time evolution of the initial wave packet can be calculated from exact diagonalisation using

$$|\psi(t)\rangle = \sum_n |n\rangle e^{-iE_n t} \langle n|\psi(0)\rangle, \tag{6.25}$$

where $|n\rangle$ are the eigenstates of the Hamiltonian with energy E_n; given $|\psi(t)\rangle$ in the site basis, μ_p follows straightforwardly. As the details of the expansion dynamics depend on the choice of $|\psi(0)\rangle$ [354], μ_p was averaged over all initial sites. Figure 6.7(a) shows the

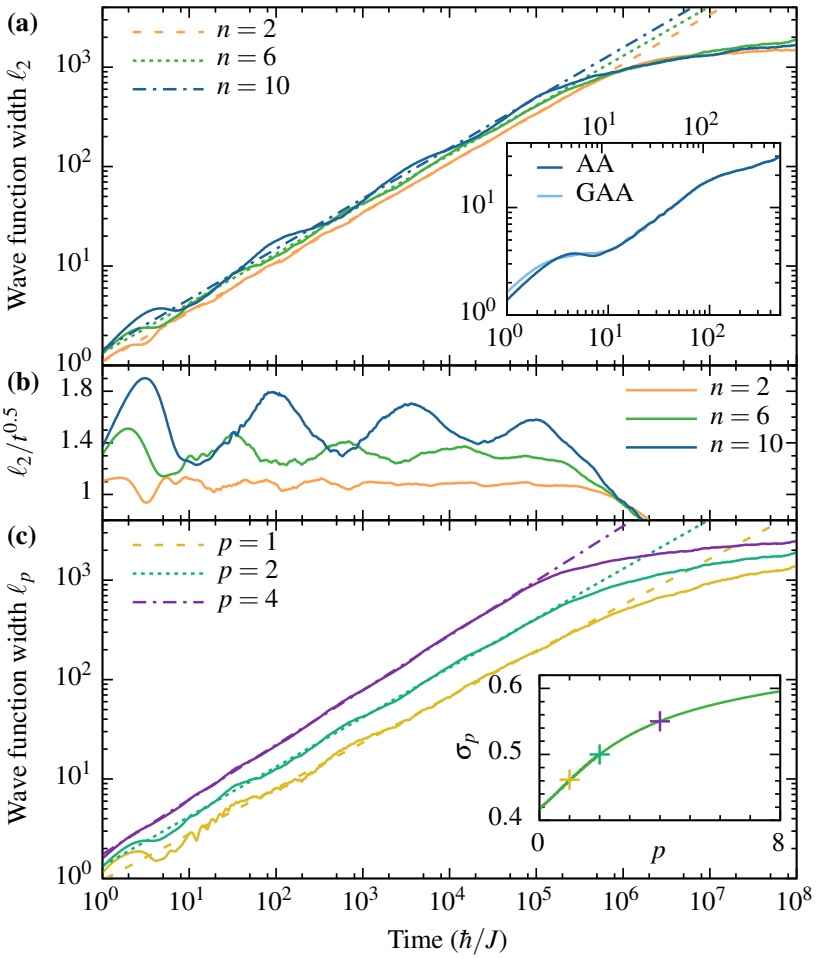

Figure 6.7. **(a)** RMS wave function width ℓ_2 in the critical Aubry–André model for rational approximations of $\beta = [0;\overline{n}]$ ($n = 2, 6, 10$) as a function of time for a state initially localised on a single site, averaged over the initial site (solid lines). For all values of β, the expansion is described well by the power law $\ell_2 \propto t^{1/2}$ (dashed and dotted lines). Convergence to a constant value at long times is a finite-size effect.
Inset: comparison of $\ell_2(t)$ for the simple ($\lambda = 2$: AA, bottom time axis) and generalised [$(\lambda, \mu, \phi) = (2, 1/2, 0)$: GAA, top time axis] Aubry–André models with $\beta = [0;\overline{10}]$. Except for very short times, the two curves are related by time dilation: $\ell_2^{\mathrm{GAA}}(t) = \ell_2^{\mathrm{AA}}(0.867t)$.
(b) $\ell_2/t^{1/2}$ for the same expansions. As before, the ratio tends to a constant at long times. The log-periodic oscillations are a consequence of the discrete RG steps admitted by the quasiperiodic Hamiltonian [329–333]. Since increasing time scales probe increasingly broad ranges of length scales in different parts of the spectrum, these oscillations are washed out at long times.
(c) ℓ_p for $\beta = [0;\overline{6}]$ and $p = 1, 2, 4$ in the same setup. For each p, ℓ_p increases as a power law; however, the critical exponents σ_p depend on p [$\sigma_1 = 0.4616(12)$, $\sigma_2 = 1/2$, $\sigma_4 = 0.5500(4)$].
Inset: comparison of σ_p calculated from the multifractal spectrum using (6.31) (solid line) to the exponents obtained numerically (coloured crosses).
Figure taken from Ref. 5.

evolution of the RMS width ℓ_2 for periodic approximations of $\beta = [0; \overline{n}]$ ($n = 2, 6, 10$) in the critical Aubry–André model. Apart from finite-size effects, each expansion follows an approximate power law: fitting one to each plot results in the same diffusion exponent $\sigma \approx 0.5$ within fitting error. Similar behaviour has previously been found for other values of β, too [355].

On the other hand, σ for a fixed value of β does depend on p, as shown in Fig. 6.7(c) for $\beta = [0; \overline{6}]$ and $p = 1, 2, 4$. Such an anomalous diffusion behaviour has been observed for the Fibonacci quasicrystal as well [356–358], but not outside quasiperiodic systems.

Beside the Aubry–André case, $\ell_p(t)$ was calculated using the same method for the critical point ($\lambda = 2$, $\mu = 1/2$, $\phi = 0$) of the generalised AA model (6.19). The RMS width ℓ_2 for $\beta = [0; \overline{10}]$ is plotted in the inset of Fig. 6.7(a) together with that of the simple AA model. The exponents of approximate power laws match, together with the structure of oscillations around it:

$$\ell_2^{\mathrm{GAA}}(t) = \ell_2^{\mathrm{AA}}(0.867t)$$

holds accurately for all but the shortest time scales. This strongly suggests that the entire spectrum of the critical generalised Aubry–André Hamiltonian renormalises onto the $\lambda = 2$ AA model uniformly, similar to the ground state (Fig. 6.6).

Multifractal analysis

The dynamics of a quantum system is determined by the spacing between its energy levels; therefore, an appropriate description of the spectrum should capture the dynamical behaviour discussed above. To elucidate this connection, consider the critical Aubry–André model with a quadratic irrational β.[†] The only natural length and time scales of the problem are the lattice spacing a_0 and the "hopping time" $\hbar/J$; μ_p scales with these as

$$\mu_p(t; x_0, \beta) = a_0^p m_p(Jt; x_0, \beta), \tag{6.26}$$

where m_p is now a dimensionless function of dimensionless variables. Similar to numerics (Fig. 6.7), we aggregate the behaviour for different initial sites by averaging over x_0:

$$\overline{m}_p(Jt; \beta) = \lim_{N\to\infty} \frac{1}{2N+1} \sum_{x_0=-N}^{N} m_p(Jt; x_0, \beta). \tag{6.27}$$

Consider now the kth step of the renormalisation process of §6.1.1: the spectrum

[†]The case of β with aperiodic continued-fraction expansions is much less straightforward than the ground-state discussion of §6.1.2, on account of the same energy (and hence time) scale corresponding to several different length scales in different parts of the spectrum, each governed by different scaling relations [5].

consists of N_k critical subbands with incommensurate ratio β_k and effective hopping terms J_i ($1 \le i \le N_k$). Provided t is much longer than the time scales corresponding to typical band gaps, interference between bands averages out, leaving

$$m_p(Jt; x_0, \beta) \simeq \sum_{i=1}^{N_k} \left|\left\langle x_0^{(i)} \middle| x_0 \right\rangle\right|^2 N_k^p \, m_p(J_i t; x_0^{(i)}, \beta_k), \tag{6.28}$$

where $\left|x_0^{(i)}\right\rangle$ is the Wannier state of subband i living (among others) on site x_0; the factor N_k^p is due to the renormalisation of lattice spacing. To average (6.28) over lattice sites, note that each renormalised band has one Wannier state per N_k lattice sites and the sum of the overlap integrals $\left|\left\langle x_0^{(i)} \middle| x_0 \right\rangle\right|^2$ over all x_0 is 1, since the $|x_0\rangle$ form a basis. As a result, the overlap integrals average to $1/N_k$ for all lattice sites, leading to

$$\overline{m}_p(Jt; \beta) \simeq N_k^{p-1} \sum_{i=1}^{N_k} \overline{m}_p(J_i t; \beta_k). \tag{6.29}$$

Finally, take those k that correspond to full periods of the continued-fraction expansion, that is, $\beta_k = \beta$. Assuming that the expansion is governed by a power law at long times,

$$\overline{m}_p(Jt; \beta) \propto (Jt)^{p\sigma_p} \quad (Jt \to \infty),$$

Eq. (6.29) gives

$$N_k^{p-1} \sum_{i=1}^{N_k} (J_i t)^{p\sigma_p} \simeq (Jt)^{p\sigma_p} \implies \sum_{i=1}^{N_k} J_i^{p\sigma_p} \propto N_k^{1-p}. \tag{6.30}$$

Since β is a quadratic irrational, the renormalisation steps of §6.1.1 are periodic, that is, the spectrum is self-similar across a period of the continued-fraction expansion. For such self-similar spectra, a continuum of *multifractal dimensions* can be defined in a number of ways, each of which contain full information about the structure of the spectrum [5, 309, 337, 359, 360]. In terms of these dimensions, the diffusion exponent σ_p can be expressed from (6.30) as

$$\sigma_p = -\frac{\tau_{1-p}}{p}. \tag{6.31}$$

As advertised, σ does depend on p, is fully determined by the self-similarity of the spectrum, and is not equal to the inverse of the ground-state dynamical exponent z. We also found that the multifractal dimensions of the simple and generalised Aubry–André models are identical: this is required by their equivalent dynamical behaviour [inset of Fig. 6.7(a)].

Note, finally, that $\ell_2 \propto \sqrt{t}$ for all values of β considered [Fig. 6.7(a,b)]. Indeed, there is strong numerical and analytical evidence [336, 361–363] that the sum of bandwidths in the Aubry–André Hamiltonian with rational $\beta = M/N$ scales as

$$\lim_{N\to\infty} N \sum_{i=1}^{N} \Delta_i \approx 9.3299 \implies \sum_{i=1}^{N} J_i \propto N^{-1}. \tag{6.32}$$

Comparing this with (6.30) yields $\sigma_2 = 1/2$, as expected. Unlike diffusive systems, however, $\sigma_p \neq 1/2$ in general, as the underlying process is fundamentally different from a random walk.

6.2 Two-dimensional Aubry–André models: mixed spectra and partially extended states

Because strait is the gate, and narrow is the way, which leadeth unto life, and few there be that find it.

— Matthew 7:14

In this section, I consider a two-dimensional generalisation of the Aubry–André model, defined on the square lattice by the Hamiltonian

$$\begin{aligned} H = & -J \sum_{nm} \left[a^\dagger_{nm} (a_{n+1,m} + a_{n,m+1}) + \text{H.c.} \right] \\ & - J\lambda \sum_{nm} \left[\cos\{2\pi\beta(n+m)\} + \cos\{2\pi\beta(n-m)\} \right] a^\dagger_{nm} a_{nm}, \end{aligned} \tag{6.33}$$

which I shall call the *2D Aubry–André (2DAA)* model [6]. Similar to the 1D AA model, it is quasiperiodic for irrational β and admits an Aubry-type duality transformation: in terms of the momentum-space operators

$$b_{jk} = \frac{1}{\sqrt{N}} \sum_{nm} \exp\left[2\pi i \beta \{ j(n+m) + k(n-m) \} \right] a_{nm}, \tag{6.34}$$

Eq. (6.33) turns into another 2DAA model, with modulation amplitude changed from λ to $4/\lambda$.

This form of quasiperiodic modulation can readily be incorporated into existing optical lattice experiments by adding two weak 1D lattices at 45° to the main lattice axes. In all numerical studies, we set $\beta = 1/\sqrt{2}$, which amounts to setting the wave vectors of the principal lattice and the perturbation equal [298]. Crucially, Eq. (6.33) is non-separable and hence fundamentally different from earlier, separable models, where the

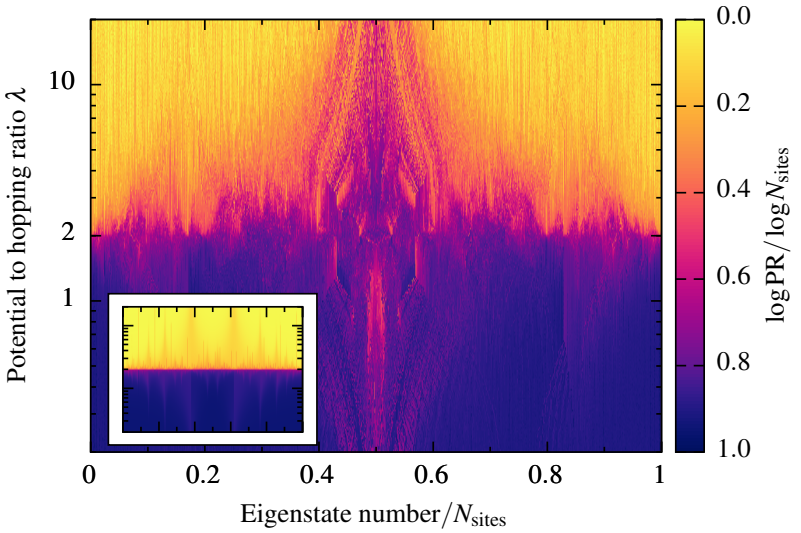

Figure 6.8. Participation ratios of all eigenstates of the 2DAA model as a function of λ for $\beta = 70/99$ ($N = 9801$ sites). A large number of partially extended states persist at $\lambda \gg 2$ in the middle of the spectrum; due to Aubry duality, a similar number of eigenstates is not fully extended for $\lambda \ll 2$.
Inset: participation ratios for the 1D AA model with $\beta = 987/1597$ ($N = 1597$). A sharp localisation transition occurs at the self-dual point $\lambda = 2$.
Figure taken from Ref. 6. Perceptionally uniform colour map chosen following Ref. 125.

non-interacting localisation transition was directly controlled by the underlying 1D AA Hamiltonian [295, 296, 364].

6.2.1 Partially extended states at strong modulation

We used exact diagonalisation to obtain the full single-particle spectrum of the 2DAA model with periodic boundary conditions, using rational approximations to $\beta = 1/\sqrt{2}$ derived from its continued-fraction expansion [346]. For each normalised eigenstate $|\psi\rangle$, we evaluated its participation ratio (PR), defined as

$$\mathrm{PR} = \left(\sum_{nm} |\psi_{nm}|^4 \right)^{-1}; \tag{6.35}$$

for a wave function evenly distributed on k sites, $\mathrm{PR} = k$. PR as a function of λ and position in the spectrum is plotted in Fig. 6.8, together with the equivalent results for the 1D AA model (inset). In the latter, one can clearly see a localisation transition in all eigenstates at $\lambda = 2$, with PR close to either 1 or the number N of sites everywhere except for a narrow region at $\lambda \approx 2$. The phase diagram of the 2D model is much more complex: there is no sharp transition at the self-dual point, but localised and delocalised

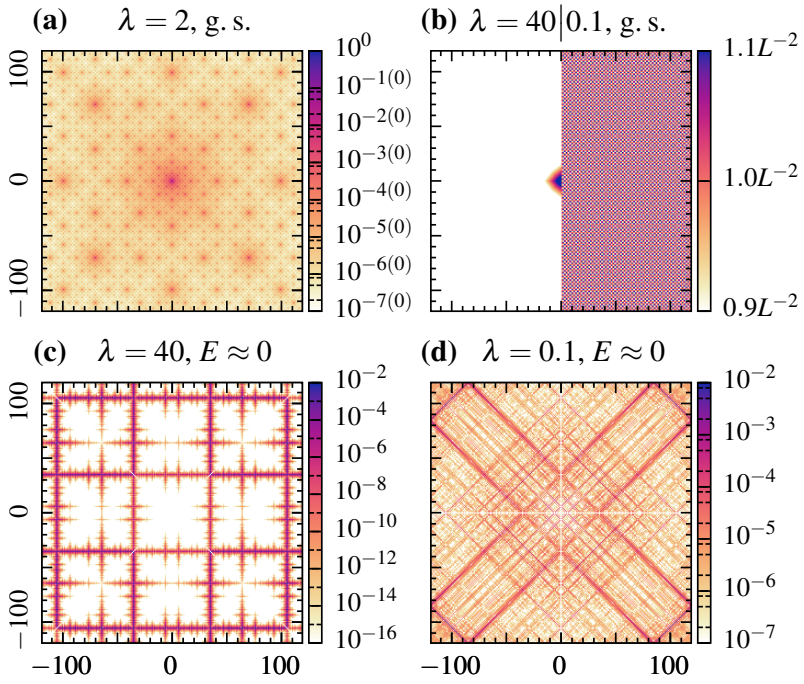

Figure 6.9. Wave function weight $|\psi_{nm}|^2$ for selected eigenstates of the 2DAA model with $\beta = 169/239$ ($L = 239$).
(a) Ground state at the self-dual point $\lambda = 2$ (log scale from 10^{-7} to 10^{0}) with a multifractal structure similar to critical eigenstates of the AA model.
(b) Ground state at $\lambda = 40$ (left half, log scale from 10^{-70} to 10^{0}) and at $\lambda = 0.1$ (linear scale to the right). The former is strongly localised with a very small localisation length ($\xi \approx 0.1$); the latter is fully extended.
(c) Median energy state ($E \approx 0$) at $\lambda = 40$. The wave function is concentrated on a few lattice lines along which the quasiperiodic potential is weaker than on average.
(d) Median energy state ($E \approx 0$) at $\lambda = 0.1$. The wave function is concentrated on a few diagonal lines of the lattice that are Aubry duals of those in (c).
Figure taken from Ref. 6. Perceptionally uniform colour map chosen following Ref. 125.

states coexist in a wide region around it. Most notably, there is a "funnel" of partially extended (PR $\sim N^{0.5}$) eigenstates in the middle of the spectrum, mixed with either localised ($\lambda \gg 2$) or extended ($\lambda \ll 2$) states. The partially extended states are not separated from localised/extended ones by a mobility edge, as expected on general grounds [334]; instead, they are interspersed in the spectrum.

To illustrate the origin of these features, we plot several representative eigenstates of the 2DAA model in Fig. 6.9. The ground states (upper panels) follow a similar pattern to the 1D case: they show fractal properties at the self-dual point and are extended and exponentially localised on either side of it. In the middle of the spectrum, however, the picture away from $\lambda = 2$ is very different (lower panels). For $\lambda \gg 2$, there exist partially

extended states, where most of the wave-function weight is concentrated on a small number of horizontal and vertical lines, with small, exponentially decaying weight close to them. Indeed, all eigenstates with significant PR follow this pattern and populate the same set of lines. For $\lambda \ll 2$, we see a similar, although less sharp, pattern concentrated on a few diagonal lines: these states are duals of the partially extended states at $\lambda \gg 2$ under the Aubry duality transformation (6.34) that includes both a Fourier transform and a 45° rotation.

The origin of these partially extended eigenstates at large λ can be explained by rewriting the potential term of (6.33) as

$$\begin{aligned} V_{nm} &= \lambda J\big[\cos\{2\pi\beta(n+m)\} + \cos\{2\pi\beta(n-m)\}\big] \\ &= 2\lambda J\cos(2\pi\beta n)\cos(2\pi\beta m) \equiv \tilde{\lambda}_n J\cos(2\pi\beta m). \end{aligned} \tag{6.36}$$

Since the effective disorder amplitudes $\tilde{\lambda}_n = 2\lambda\cos(2\pi\beta n)$ along lattice lines form a quasiperiodic sequence, there will always be lines for which $\tilde{\lambda} \ll \lambda$, that is, the disorder is much weaker than it typically is across the system. If one removed the horizontal hopping terms from (6.33), the remaining model would consist of independent AA chains with parameter $\tilde{\lambda}_n$: along lines where $|\tilde{\lambda}_n| < 2$, all eigenstates would be extended [308]. Reintroducing the horizontal hopping terms would then hybridise the AA eigenstates on different lines. By the same argument, horizontal lines also form AA models, some of which are in the extended phase: these hybridise with the aforementioned vertically extended states, leading to mesh-like eigenstates living on the quasiperiodic grid of low-disorder lines [Fig. 6.9(c)]. Since a finite fraction [approximately $2/(\pi\lambda)$] of all lines is extended, the PR of these states will, for sufficiently large system sizes, scale as L^2, i.e., as extended states. On intermediate scales, however, they appear one-dimensional, with their 2D character limited to the intersections of horizontally and vertically extended lines. Therefore, we call them *partially extended*.

On the other hand, the exponential localisation of lines with $|\tilde{\lambda}| > 2$ suppresses the statistical weight of partially extended states away from the low-disorder lines. As a result, partially extended and localised states at similar energies do not hybridise. This prevents clean mobility edges [334] and instead gives rise to the observed mixed spectrum, similar to other quasiperiodic systems [365–369].

It is indeed easy to verify that the state shown in Fig. 6.9(c) is extended precisely along the lines with $|\tilde{\lambda}| < 2$ and that for $|\tilde{\lambda}| > 2$, the effective localisation length decreases with growing local amplitude (Table 6.1). Furthermore, since the partially extended eigenstates are effectively superpositions of extended Aubry–André eigenstates, we expect their energies to lie within the spectrum of the relevant AA model. In partic-

n	±35	±105	±64	±6	±76	±93	±23	±47
$\lvert\tilde{\lambda}_n\rvert$	0.526	1.577	2.628	3.679	4.729	5.779	6.827	7.874
ξ_{1D}	——	——	3.662	1.641	1.162	0.942	0.815	0.730

Table 6.1. The eight smallest local disorder amplitudes $\tilde{\lambda}$ for the 2DAA model with $\beta = 169/239$ and $\lambda = 40$. $|\tilde{\lambda}| < 2$ for the first two: the resulting extended AA eigenstates account for the bulk of the statistical weight in Fig. 6.9(c). Subsequent lines define localised AA models, which appear in the same state as progressively shorter barbs [comparable to the AA localisation length (6.12)] close to the extended lines.

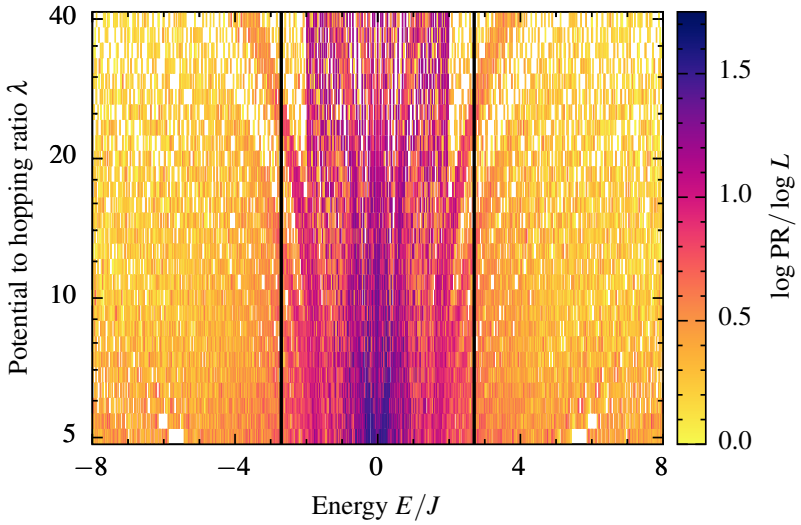

Figure 6.10. Participation ratios of the eigenstates of the 2DAA model close to zero energy, as a function of λ and energy, for $\beta = 169/239$ ($L = 239$). Regardless of the value of λ, a large number of eigenstates with energies within the bandwidth of the critical AA model (black lines) have large ($L^1..L^{1.7}$) PR; no such states occur outside of this energy window. Figure taken from Ref. 6. Perceptionally uniform colour map chosen following Ref. 125.

ular, they should only occur between the lowest and highest eigenvalues of the critical ($\lambda = 2$) AA Hamiltonian, as these bound the spectrum of the same for all $\lambda < 2$.[†] We show that this is indeed the case in Fig. 6.10: the PR of eigenstates outside of the critical bandwidth is substantially smaller than the large values attained by many eigenstates inside. Note, however, that some states remain localised even inside this energy range: these appear on lattice sites with small $|V_{nm}|$ away from the extended lines. The result is the peculiar mixed spectrum of interspersed localised and partially extended states.

[†]For $\beta = 1/\sqrt{2}$, we find these bounds by exact diagonalisation to be $\approx \pm 2.703J$ (black lines in Fig. 6.10).

6.2.2 Quench dynamics

In experimental settings, dynamics following a quantum quench from a given initial state is often easier to access than individual eigenstates. Nevertheless, the coexistence of localised and partially extended eigenstates at $\lambda \gg 2$ and the fact that the latter are largely confined on a special set of low-disorder lattice lines indicate that the same lines will also show peculiar dynamical properties. To confirm this intuition, we have considered the expansion dynamics of states initially localised on a single lattice site (n, m) under the 2DAA Hamiltonian [cf. (6.25)]. In particular, we focus on the diagonal ensemble

$$\rho_{nm} = \sum_{\psi} \left|\langle\psi|nm\rangle\right|^2 |\psi\rangle\langle\psi|, \tag{6.37}$$

the long-time average of the density matrix $|\Psi(t)\rangle\langle\Psi(t)|$, where $|\Psi(0)\rangle = |nm\rangle$ and the sum runs over the eigenstates of (6.33). Given the density distribution due to this diagonal ensemble, $\varrho_{nm}(n', m') = \langle n'm'|\rho_{nm}|n'm'\rangle$, we define its participation ratio similarly to (6.35) as

$$\mathrm{PR}_{nm} = \Big(\sum_{n'm'} \left|\langle n'm'|\rho_{nm}|n'm'\rangle\right|^2 \Big)^{-1}. \tag{6.38}$$

Broadly speaking, this participation ratio captures how many lattice sites a particle will reach starting from an initial state localised to a single site. It is plotted in Fig. 6.11 as a function of the initial site for $\lambda = 2$ and 40, along with two representative diagonal ensembles for each. At the critical point, diagonal ensembles show fractal properties but altogether appear delocalised on the simulated length scales; indeed, $\mathrm{PR} \approx L^2$ for most initial sites. For $\lambda \gg 2$, the difference between low-disorder lines and the rest of the lattice appears very pointedly in the participation ratios. For most initial sites (including the origin), the bulk of the probability distribution remains localised close to the starting point, with exponentially suppressed probability of reaching the network of low-disorder lines. On the contrary, starting from such a line leads to fast delocalisation across this network, in a pattern similar to partially extended eigenstates [Fig. 6.9(c)]. By evaluating $|\Psi(t)\rangle$ explicitly (§6.1.3), it can be shown that the expansion dynamics along these lines occurs at constant speed, i.e., it is ballistic [6].

To quantify whether the diagonal ensembles in Fig. 6.11 are extended, we evaluated the total radial probability $P(r)$ and mean radial density $\overline{\rho(r)}$, defined by

$$P(r) = \sum_{d=r-1/2}^{r+1/2} \langle n'm'|\rho_{nm}|n'm'\rangle, \qquad \overline{\rho(r)} = \frac{P(r)}{N(r)}, \tag{6.39}$$

where $N(r)$ is the number of sites whose distance $d = \sqrt{(m-m')^2 + (n-n')^2}$ from the

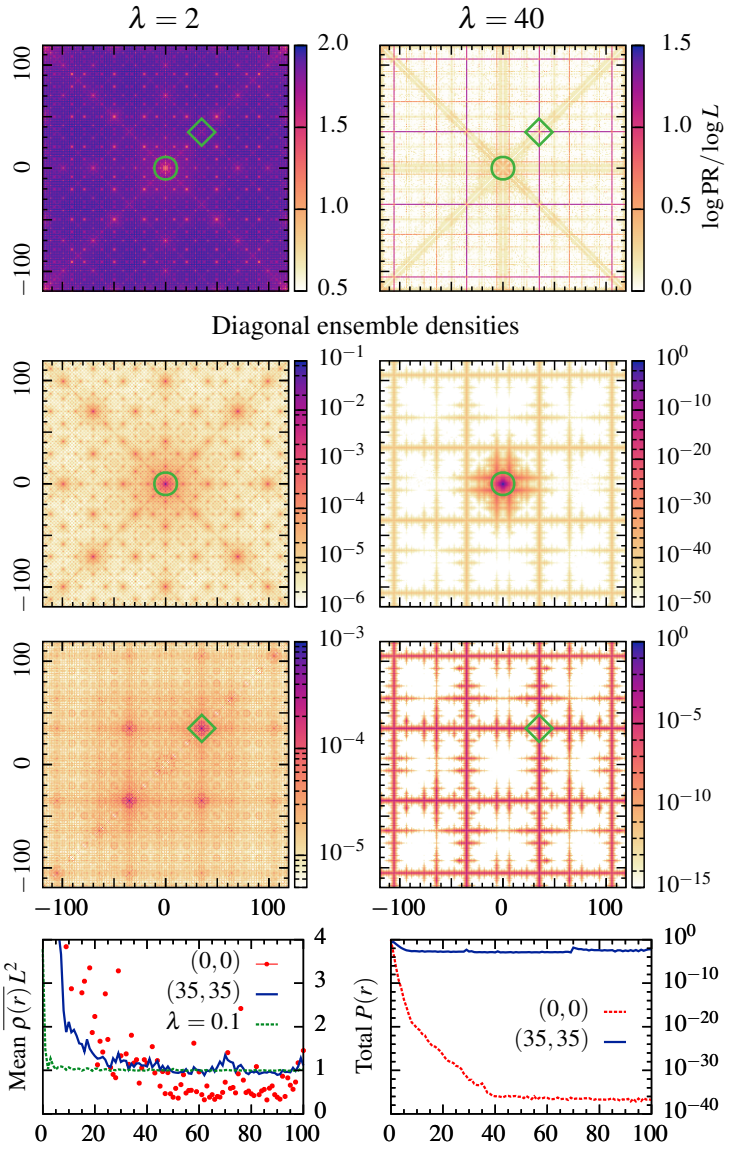

Figure 6.11. **Top panels**: Diagonal-ensemble PR (6.38) as a function of the initial site for $\lambda = 2$ and 40. At the self-dual point, the distribution appears almost uniformly delocalised for most initial sites; at $\lambda \gg 2$, by contrast, all sites except those along low-disorder lines show strongly localised dynamics.
Middle panels: Diagonal-ensemble densities $\varrho_{nm}(n', m')$ for initial sites $(n, m) = (0, 0)$ and $(35, 35)$ (green symbols) for $\lambda = 2, 40$. At the self-dual point, special initial sites like the origin lead to pronounced fractal dynamics; for more generic sites, the late-time distribution is almost uniform. At strong modulation, delocalised dynamics occurs only on low-disorder lines.
Bottom panels: The mean radial density $\rho(r)$ tends to a constant $\approx L^{-2}$ for almost all initial sites for $\lambda \leq 2$, indicating 2D extended dynamics [left; blue and red: $\lambda = 2$; green: $\lambda = 0.1$, starting from $(35, 35)$]. For $\lambda \gg 2$ (right), the total radial probability $P(r)$ is approximately constant for initial sites on low-disorder lines, indicating partially extended (quasi-1D) dynamics (blue curve); for a generic initial site (red curve), exponential decay is capped at very low levels by the contribution of delocalised lines.
Figure taken from Ref. 6. Perceptionally uniform colour map chosen following Ref. 125.

initial site is between $r - 1/2$ and $r + 1/2$. A constant $P(r)$ indicates states that are uniformly extended along individual lines, while a constant $\overline{\rho(r)}$ corresponds to states extended over the whole plane. Indeed, for $\lambda \leq 2$, $\overline{\rho(r)}$ tends to a constant $\approx L^{-2}$ for most initial sites, indicating almost uniform distributions compatible with extended states. For $\lambda \gg 2$, $P(r)$ decays exponentially for most initial sites before it is eventually capped by the contribution of delocalised lines. Starting from one of these lines, by contrast, leads to a nearly constant $P(r)$, consistent with uniform delocalisation in one dimension.

6.2.3 Generality of effect

The 2DAA model (6.33) is fine-tuned in the sense that the quasiperiodic modulation must be applied at precisely $\pm 45^\circ$ to get infinitely long weakly modulated lines. In Ref. 6, I studied two ways in which an experimentally realised quasicrystal may differ from this idealised Hamiltonian:

- The modulating potential may be tilted from its ideal direction. For simplicity, I focused on models in which the wave vectors of the two cosine potentials remain perpendicular, but no longer make 45° angles with the lattice axes. These models also admit an Aubry duality transformation similar to (6.34), marking out $\lambda = 2$ as a special self-dual point for all of them. For modulation angles close to 45°, the structure of the spectrum and eigenstates remains similar as long as λ is not much larger than 2. One can understand these eigenstates as still living on low-disorder lines, which, however, no longer align perfectly with the lattice axes; as long as the modulation is not too strong, it nevertheless remains possible for the wave function to hop between adjacent lattice lines and thus follow the low-disorder region. For larger rotation angles, this structure is effectively destroyed; nevertheless, there is still no sharp transition at $\lambda = 2$ for all eigenstates, as also found recently in Ref. 368. At the other extreme, if the wave vectors of the modulating potential line up with the lattice axes, the Hamiltonian becomes separable, and so its eigenstates are outer products of Aubry–André eigenstates: like the AA model, therefore, it undergoes a sharp localisation transition at the self-dual point.

- Random disorder always leads to exponential localisation in two dimensions, albeit with exponentially large localisation lengths [370]. To understand the effect of such disorder on the 2DAA model, I evaluated the diagonal-ensemble PR (6.38) for several disorder configurations on top of quasiperiodic modulation near the self-dual point. The fractal features characteristic of this point (Fig. 6.11) become less and less pronounced with increasing disorder; instead, initial sites on or near the

weak-disorder network observed in the $\lambda \gg 2$ case retain a substantially larger PR than most other lattice sites. Individual diagonal ensembles display a strong maximum near the initial site, consistent with localisation; however, the mean radial density $\overline{\rho(r)}$ only decays as a power law, rather than exponentially. This is probably so because even at the relatively large disorder strengths used in our simulations, 2D Anderson localisation has a larger correlation length than the system sizes amenable to exact diagonalisation: accordingly, certain exotic features of the 2DAA models would be amenable to experiment as well.

In future work, it will be interesting to study the effect of the delocalised "rare lines" on MBL. These act as thermalising "seeds" that enclose finite [average linear size: $O(\lambda)$] Anderson-localised bubbles: this geometry is qualitatively different from the typical notion of 2D rare regions and may give rise to novel types of Griffiths effects [371]. Studying MBL on the large systems required by this geometry is numerically challenging: semiclassical methods (such as the Gross–Pitaevskii equation) provide a useful first approximation in experimentally relevant scenarios [372]. Finally, certain 1D quasiperiodic systems (including the Aubry–André model) exhibit a form of bulk–boundary correspondence, which results in end modes topologically protected from random disorder [373]: a similar mechanism may protect the partially extended states of the 2DAA model as well.

7

Conclusion

"Begin at the beginning," the King said, very gravely,
"and go on till you come to the end: then stop."
— Lewis Carroll, *Alice in Wonderland*

Throughout this dissertation, I developed several approaches to tackle the challenging complexity of quantum-spin-liquid systems. The first of these captures QSL physics in a semiclassical large-S approximation: in this limit, quantum correlations and uncertainty relations become immaterial, making the models amenable to highly efficient classical (Landau–Lifshitz) dynamical and thermodynamic simulations (§2.2). Semiclassical techniques have previously been used successfully on a variety of models; standard approaches, however, have hitherto failed to capture spin-liquid physics in systems where it is underpinned by local constraints and higher-order perturbative processes (§2.3). This includes historically and experimentally significant models, such as quantum vertex and dimer models and quantum spin ice (QSI), where the large-S approximation yields frustrated Ising models, devoid of spin-liquid dynamics.

I found that this limitation can be eliminated by employing the large-S approximation directly on the perturbative ring-exchange Hamiltonian that underpins the QSL physics of these models. This insight allowed us to develop a systematic approach to deriving effective field theories for Rokhsar–Kivelson quantum dimer models [1]: this approach recovers the familiar height and flux mappings, as well as qualitatively correct U(1) gauge theories, on bipartite lattices (chapter 3). I complemented this approach with semiclassical numerical simulations of the ring-exchange Hamiltonian in quantum spin ice [2]: in addition to the photon modes readily captured by the field theory, these simulations allow convenient access to the elusive gapped magnetic monopole (vison) quasiparticles of QSI, their interactions and thermodynamic properties (§2.4). These approaches greatly extend the power of semiclassical methods to capture QSL physics, covering the phenomenology of most U(1) liquid phases.

An important remaining challenge of large-S approaches is capturing the dynamics of gapped excitations, such as the U(1) visons in QSI: this dynamics is due to instanton effects, which are challenging to include in large-S analytical calculations (§3.4.3), and are not captured at all by low-temperature dynamical simulations (§2.3). This issue is particularly pointed in $\mathbb{Z}_2$ spin liquids (e.g., dimer models on non-bipartite lattices), where all relevant excitations are gapped and the gauge symmetry itself is discrete, which is again at odds with the smooth evolution imposed by semiclassics. It will be interesting to see whether explicitly spin-1/2 techniques based on semiclassics, such as the truncated Wigner approximation [374] or classical stochastic processes [375–377], are able to provide a more accurate grasp of the dynamics and thermodynamics of the true spin liquid models while retaining the intuitive power of large-S approaches.

The second approach aims to find the exact ground-state wave function of QSL models by employing neural networks as wave-function ansätze. The principal appeal of this technique is the flexibility of deep neural networks, central to the success of modern machine learning, which could allow them to represent wave functions with arbitrarily complex structures in an unbiased fashion (§4.2). Contrary to these expectations, I found that variational methods with such neural quantum states (NQS) generally struggle to converge to ground states that contain both positive and negative wave-function amplitudes [3]. I designed an ansatz (§4.3.1) that was able to capture the Marshall sign rule of unfrustrated antiferromagnets on the square lattice while enforcing translation invariance, an improvement over previous approaches with deep neural networks. The more complicated sign structure in the spin-liquid phase of the J_1–J_2 Heisenberg model on the square lattice, however, remains an open challenge, both for our ansatz and for most state-of-the-art approaches (§4.4.4).

This sign problem appears to be an intrinsic limitation of all NQS ansätze that represent amplitudes and phases of the wave function independently, on account of the latter having an extremely complex structure in sign-problematic ground states apart from special cases like the Marshall sign rule [198]. This undermines the learning ability of neural networks, which depends on efficiently generalising relatively simple patterns. The key open question of NQS approaches is, therefore, whether network architectures can combine sign and amplitude information in a way that allows changing the sign of the wave function continuously during training: this is critical for sign-problematic Hamiltonians as imaginary-time evolution under them (which is emulated by the training process) leads to destructive interferences and thus frequent sign changes. Thus far, the most successful NQS ansätze for spin systems achieve this by exploiting the connection between antiferromagnets and interacting fermion systems, a piece of deep physical insight that does not generalise well to non-Heisenberg models (§4.5.4). In future work,

I aim to explore whether simpler (and thus more generally usable) ansätze are capable of capturing these complicated ground states.

Additionally, I reported on a theory–experiment collaboration on the low-temperature magnetism of the pyrochlore iridate $Ho_2Ir_2O_7$ using magnetisation and magnetoresistance measurements, as well as numerical modelling (chapter 5). We found that the resistivity of $Ho_2Ir_2O_7$ is highly and specifically sensitive to the density of emergent magnetic monopoles in spin-ice iridates (§5.3.3): I discussed scattering mechanisms underlying this behaviour, which appear equally relevant in related materials, providing a novel, sensitive probe of magnetic monopoles in spin ice. Furthermore, holmium spins, which form a dipolar spin ice coupled to the antiferromagnetically ordered iridium domains, mediate an effective coupling between external magnetic fields along the [111] crystallographic direction and the same iridium domains (§5.2.2). This allows the distribution of the two symmetry-broken domain types to be controlled by an external field, a highly unusual scenario for antiferromagnetic materials (§5.3.2). Since the mechanism requires relatively large fields, the antiferromagnetic domains remain protected from stray fields, making the setup potentially useful for antiferromagnetic spintronic technology [252]. Further research in this direction could focus on replicating a similar interplay between antiferromagnetic order and spin-liquid physics in systems with larger coupling energy scales, creating materials useful for technological applications at room temperature.

Finally, I discussed unique localisation properties of quasiperiodic Hamiltonians in one and two dimensions (chapter 6). In particular, I found that the quantum localisation transitions of the 1D Aubry–André model and several related Hamiltonians exhibit a form of universality, which, unlike conventional second-order phase transitions, is not captured by simple power laws and critical exponents [5]. Furthermore, I introduced a 2D generalisation of the Aubry–André model on the square lattice, which exhibits a peculiar mixed spectrum of partially extended and localised eigenstates for arbitrarily strong modulation [6]: the former are constrained onto a network of weakly modulated lines of the lattice, allowing them to coexist with localised states at the same energy, without mobility edges. In the context of many-body localisation, such "rare lines" may give rise to novel kinds of Griffiths effects, which may shed light on the stability of MBL in 2D in general.

A

Semiclassical simulations of quantum spin ice

A.1 Vector calculus on the pyrochlore lattice

To construct the lattice gauge theory (2.5), it is important to observe that the pyrochlore lattice is made up of link midpoints of a diamond lattice whose sites are tetrahedron centres of the former; that is, variables on pyrochlore sites can be thought of as link variables of the same diamond lattice. In order to turn these variables into a lattice vector field v of the kind described in §1.1.2, we have to make the links directed and define the field such that $v_{rr'} = -v_{r'r}$. Importantly, the diamond lattice is bipartite, that is, it can be separated into two sublattices A and B (the "up" and "down" tetrahedra) without any links within a sublattice: this allows us to turn any set of link variables x into such a vector field by setting

$$v_{rr'} = \begin{cases} x & \text{if } r \in A, r' \in B \\ -x & \text{if } r \in B, r' \in A, \end{cases} \tag{A.1}$$

analogous to (2.4); note that no similarly consistent definition of lattice vector fields can be given for non-bipartite lattices.

Lattice scalar fields can also be defined: they live on the sites of the diamond lattice and, like scalar fields in $\mathbb{R}^d$, their sign is uniquely defined. This definition allows us to introduce the lattice divergence of a vector field and the lattice gradient of a scalar field:

$$(\operatorname{grad} u)_{rr'} = u_{\vec{r}'} - u_{\vec{r}}; \qquad (\operatorname{div} a)_{\vec{r}} = \sum_{\vec{r}':\langle rr'\rangle} a_{rr'}. \tag{A.2}$$

It is easy to see that the gradient is a well-defined vector field, while div a is a scalar field. As an example, the lattice divergence of the electric field e is the spinon charge of the corresponding tetrahedron, in analogy with Gauss' law.

The lattice curl is defined on the plaquettes of the lattice, by summing a vector field around the links of each plaquette. For the concrete example of the magnetic field $b = \operatorname{curl} a$,

$$b \equiv \operatorname{curl} a = \sum_{i=1}^{6} a_{r_i r_{i+1}} = \phi_1 - \phi_2 + \phi_3 - \phi_4 + \phi_5 - \phi_6; \tag{A.3}$$

the alternating signs in the final form appear because the tetrahedra around a plaquette belong to alternating diamond sublattices. It is natural to define the lattice curl b to live on the centres of the plaquettes, which in turn form a *dual* pyrochlore lattice: this and the fact that the curl in standard vector calculus is a vector field motivates making b a lattice vector field on the links of the corresponding dual diamond lattice [cf. Fig. 1.6(b)]. Conveniently, the definition (A.3) has a sign ambiguity: reversing the direction in which the plaquette is traversed flips the signs of all a fields, and thus of b itself. Therefore, b on a *directed* link of the dual diamond lattice can be defined by (A.3), using the right-hand rule to orient the plaquette, analogous to the right-hand rule in standard vector calculus. It follows easily from the definitions (A.2, A.3) that $\operatorname{curl}\operatorname{grad} u = 0$ for any scalar field u; with some effort, it can also be shown that $\operatorname{div}\operatorname{curl} a = 0$ (this object lives on the dual diamond lattice); that is, the operators grad, div, and curl behave like their counterparts in continuum vector calculus.

The lattice gradient and divergence defined by (A.2) remain valid in arbitrary dimensions, for any (bipartite) lattice. Like its continuum counterpart, the lattice curl is only well-defined as a vector field in 3D, including lattices other than the pyrochlore [see the discussion of Ref. 15 of U(1) gauge theories on a cubic lattice]. In 2D, the definition (A.3) can still be evaluated, but as the plaquettes cannot be regarded as links of a dual lattice anymore, it is no longer a vector field, so objects like $\operatorname{div}\operatorname{curl} a$ become ill-defined.

A.2 Quadratic estimates of the vison cost and interaction

Following Ref. 17, we estimate the zero-temperature gap and interaction strength of two visons in a quadratic theory, where the energy of the magnetic field on each plaquette, $-g\cos b$, is approximated as $-g + gb^2/2$. Let the number of visons on each site of the dual diamond lattice be $n_{\vec{r}}$; we want to find the lowest-energy configuration of the magnetic field for which $\operatorname{div} b = 2\pi n_{\vec{r}}$. To do so, we consider a lattice version of Helmholtz's theorem: any vector field $b_{rr'}$ on the dual pyrochlore lattice with periodic boundary conditions can be written as

$$b = -\operatorname{grad} u + \operatorname{curl} a + C^{\mu}, \tag{A.4}$$

where u is a scalar potential defined on the dual diamond lattice, a is a vector potential defined on the original pyrochlore lattice (both with the same periodic boundary conditions as b), and C^μ is a background field that only depends on the four FCC sublattices $\mu = 0, 1, 2, 3$ of the dual pyrochlore lattice. Now, the cross-terms between the three components vanish in the sum of b^2 over all sites, that is, the approximate total energy of a field configuration can be written as

$$E - E_0 \simeq \frac{g}{2} \sum_{\langle rr' \rangle} b^2 = \frac{g}{2} \sum_{\langle rr' \rangle} \left[(\operatorname{grad} u)^2 + (\operatorname{curl} a)^2 + (C^\mu)^2 \right]. \tag{A.5}$$

Since u is uniquely determined by n, the energy is minimised if curl a is identically zero, that is, there are no non-gradient components of b. From here on, we also neglect the C^μ term in (A.4): these give rise to surface terms proportional to the overall "magnetisation" of the configuration within periodic boundary conditions [2, 127].

The scalar potential u can be obtained by requiring that

$$\operatorname{div} b = -\operatorname{div} \operatorname{grad} u = 2\pi n. \tag{A.6}$$

The divergence and the gradient can be expressed in reciprocal space as

$$(\operatorname{grad} u)^\mu(\vec{q}) = -\left[\tilde{M}^\dagger(\vec{q})\right]^{\mu\alpha} u^\alpha(\vec{q}) \tag{A.7}$$

$$(\operatorname{div} b)^\alpha(\vec{q}) = \left[\tilde{M}(\vec{q})\right]^{\alpha\mu} b^\mu(\vec{q}), \tag{A.8}$$

where the indices $\alpha = \pm$ and $\mu = 0, 1, 2, 3$ indicate the FCC sublattices of the diamond and pyrochlore lattices, respectively, and

$$\tilde{M}(\vec{q}) = \begin{pmatrix} e^{i\vec{q}\cdot\vec{r}_0/2} & e^{i\vec{q}\cdot\vec{r}_1/2} & e^{i\vec{q}\cdot\vec{r}_2/2} & e^{i\vec{q}\cdot\vec{r}_3/2} \\ -e^{-i\vec{q}\cdot\vec{r}_0/2} & -e^{-i\vec{q}\cdot\vec{r}_1/2} & -e^{-i\vec{q}\cdot\vec{r}_2/2} & -e^{-i\vec{q}\cdot\vec{r}_3/2} \end{pmatrix}, \tag{A.9}$$

where $\vec{r}_0 = a_0[111]/4$, $\vec{r}_1 = a_0[1\bar{1}\bar{1}]/4$, $\vec{r}_2 = a_0[\bar{1}1\bar{1}]/4$, and $\vec{r}_3 = a_0[\bar{1}\bar{1}1]/4$ are the vectors pointing from a "+" diamond lattice site to its nearest neighbours. It follows that (A.6) can be written as (we suppress the arguments $\vec{q}$ from now on)

$$\tilde{M}\tilde{M}^\dagger \tilde{u} = 2\pi\tilde{n} \implies \tilde{b} = \tilde{M}^\dagger \tilde{u} = 2\pi \tilde{M}^\dagger \left(\tilde{M}\tilde{M}^\dagger\right)^{-1} \tilde{n},$$

where we introduced the vectors $\tilde{n} = [n^+, n^-]$, $\tilde{u} = [u^+, u^-]$, and $\tilde{b} = [b^0, b^1, b^2, b^3]$ for convenience. Equation (A.10) gives all $\vec{q} \neq 0$ Fourier components of the ground state magnetic field configuration.[†] Now, the total energy of this configuration in the

[†] $\tilde{M}\tilde{M}^\dagger$ is singular at the Γ point: this is the origin of the constant term in (A.4).

quadratic approximation is

$$E - E_0 = \frac{g}{2}\sum_{\vec{r}} b^2(\vec{r}) = \frac{g}{2N}\sum_{\mu,\vec{q}} \left|b^\mu(\vec{q})\right|^2 \simeq \frac{(2\pi)^2 g}{2N}\sum_{\vec{q}\neq 0} \tilde{n}^\dagger \left(\tilde{M}\tilde{M}^\dagger\right)^{-1} \tilde{M}\tilde{M}^\dagger \left(\tilde{M}\tilde{M}^\dagger\right)^{-1}\tilde{n}$$
$$\longrightarrow 2\pi^2 g V_{\text{cell}} \int_{\text{BZ}} \frac{\mathrm{d}^3 q}{(2\pi)^3} \tilde{n}^\dagger \left(\tilde{M}\tilde{M}^\dagger\right)^{-1}\tilde{n}, \tag{A.10}$$

where N is the number of FCC unit cells of volume $V_{\text{cell}} = a_0^3/4$, and (cf. Ref. 12)

$$\tilde{M}M^\dagger = 4\begin{pmatrix} 1 & -\gamma \\ -\gamma^* & 1 \end{pmatrix} \implies \left(\tilde{M}\tilde{M}^\dagger\right)^{-1} = \frac{1}{4(1-|\gamma|^2)}\begin{pmatrix} 1 & \gamma \\ \gamma^* & 1 \end{pmatrix}. \tag{A.11}$$
$$\left[\gamma = \frac{1}{4}\sum e^{i\vec{q}\cdot\vec{r}_\mu} = \cos(q_x/4)\cos(q_y/4)\cos(q_z/4) - i\sin(q_x/4)\sin(q_y/4)\sin(q_z/4)\right]$$

To obtain the interaction energy of two visons from the general form (A.10), consider a positive and a negative vison, both on the "+" sublattice, distance $\vec{R}$ apart. That is, let $n(\vec{0}) = +1$, $n(\vec{R}) = -1$, and $n(\vec{r}) = 0$ otherwise. In reciprocal space, this yields $n^+(\vec{q}) = 1 - e^{i\vec{q}\cdot\vec{R}}$ and $n^-(\vec{q}) = 0$; substituting this into (A.10) gives

$$E - E_0 = \frac{\pi^2 g V_{\text{cell}}}{2}\int_{\text{BZ}} \frac{\mathrm{d}^3 q}{(2\pi)^3}\frac{|1 - e^{i\vec{q}\cdot\vec{R}}|^2}{1-|\gamma^2|}$$
$$= 2\times\frac{g a_0^3\pi^2}{8}\int_{\text{BZ}}\frac{\mathrm{d}^3 q}{(2\pi)^3}\frac{1}{1-|\gamma|^2} - \frac{g a_0^3\pi^2}{4}\int_{\text{BZ}}\frac{\mathrm{d}^3 q}{(2\pi)^3}\frac{\cos(\vec{q}\cdot\vec{R})}{1-|\gamma|^2}. \tag{A.12}$$

The two terms of (A.12) can be regarded as the energies of the two isolated visons and their interaction, respectively. The bare vison energy works out to be $\mu \approx 8.848g$. For large $\vec{R}$, the second term integrates to zero over most of the Brillouin zone due to the rapidly oscillating cosine factor; the only exception is the vicinity of the Γ point, where $\gamma = 1$, so the integrand diverges. There, $|\gamma|^2 \approx 1 - q^2 a_0^2/16$, and so the integral becomes

$$E_{\text{int.}} \approx -4 g a_0 \pi^2 \int \frac{\mathrm{d}^3 q}{(2\pi)^3}\frac{\cos(\vec{q}\cdot\vec{R})}{q^2} = -\frac{g a_0 \pi}{R}, \tag{A.13}$$

which demonstrates the effective Coulomb interaction between visons at large distances, arising from a purely short-range Hamiltonian.

A.3 Pinch-point blurring due to free and bound visons

We derive the contribution of visons to the equal-time correlator of the emergent magnetic field, $\langle b(\vec{q})b(-\vec{q})\rangle$, for several arrangements of visons that are realised in the

semiclassical quantum-spin-ice model. We assume throughout that there is no interaction between photons and visons, that is, the overall correlator can be written as the sum of independent vison and photon contributions.

Let the vison number on each site of the dual diamond lattice be $n_{\vec{r}}$, which can be Fourier transformed into $n^\alpha(\vec{q})$, as done in Appendix A.2. The contribution of these visons to the magnetic field is given by (A.10):

$$b(\vec{q}) = \sum_\mu b^\mu(\vec{q}) = 2\pi \mathbb{1}\tilde{M}^\dagger\left(\tilde{M}\tilde{M}^\dagger\right)^{-1}\tilde{n}(\vec{q}) = \tilde{v}\tilde{n}(\vec{q}), \tag{A.14}$$

where $\mathbb{1}$ is the row four-vector all entries of which are 1, and we introduce the row two-vector

$$\tilde{v} = 2\pi \mathbb{1}\tilde{M}^\dagger(\tilde{M}\tilde{M}^\dagger)^{-1} = \frac{2\pi}{1-|\gamma|^2}\left[\kappa^* - \gamma^*\kappa,\ \ \gamma\kappa^* - \kappa\right], \tag{A.15}$$

where $\kappa = \frac{1}{4}\sum_\mu e^{i\vec{q}\cdot\vec{r}_\mu/2}$. The vison contribution to the correlator of $b(\vec{q})$ follows from that of $n(\vec{q})$ as

$$\langle b(-\vec{q})b(\vec{q})\rangle_{\text{vison}} = \sum_{\alpha,\beta}\overline{v^\alpha}v^\beta\left\langle n^\alpha(-\vec{q})n^\beta(\vec{q})\right\rangle. \tag{A.16}$$

In the following, I derive the correlator of vison number, and hence the vison contribution to the magnetic-field correlators, for vison arrangements relevant to our model. For convenience, we focus on the behaviour along the q_z axis, where the photon pinch-point contribution vanishes [12], so the $\langle bb\rangle$ correlator measured along it is entirely due to visons and allows for direct comparison with the result in Eq. (A.16).

A.3.1 Debye plasma of dissociated visons

Let the visons interact through the reduced Coulomb interaction $\beta V(r) = Kq_iq_j/(4\pi r)$. Assume that the density ρ of visons is small, that is, both their typical separation and the Debye screening length of the resulting plasma is much larger than the lattice spacing. Accordingly, we focus on the vicinity of Γ points in reciprocal space. From standard Debye–Hückel theory [64], the pair correlation function of visons is

$$g_{ab}(r) = e^{-\beta q_a q_b \phi(r)} \approx 1 - \beta q_a q_b \phi(r); \qquad \beta\phi(r) = \frac{Ke^{-r/\xi_D}}{4\pi r}, \tag{A.17}$$

where $\xi_D = (K\rho)^{-1/2}$ is the Debye screening length and $a, b = \pm$ denote the positive and negative vison species. To go from (A.17) to the correlators $\langle n^\alpha n^\beta\rangle$, we note that the long-wavelength theory does not discriminate between the two sublattices. Therefore, the density of visons in each is $\rho/2$ and $\langle n^\alpha n^\beta\rangle$ does not depend on α and β in the long-

wavelength limit. In this approximation,

$$\left\langle n^{\alpha}(\vec{R})n^{\beta}(\vec{R}+\vec{r})\right\rangle \propto \left(\frac{\rho}{4}\right)^{2}\sum_{a,b=\pm} q_a q_b g_{ab}(r) = -\frac{\rho^2}{4}\frac{Ke^{-r/\xi_{\mathrm{D}}}}{4\pi r}$$
$$\left\langle n^{\alpha}(-\vec{q})n^{\beta}(\vec{q})\right\rangle_2 = -\frac{K\rho^2}{4}\frac{1}{\vec{q}^2+\xi_{\mathrm{D}}^{-2}} = -\frac{\rho}{4}\frac{\xi_{\mathrm{D}}^{-2}}{\vec{q}^2+\xi_{\mathrm{D}}^{-2}}. \tag{A.18}$$

The derivation above captures the correlation between pairs of two visons; however, the correlation of visons with themselves also contributes to the $\langle nn\rangle$ correlator. This contribution is clearly a δ-function in real space, so it only couples each sublattice to itself; after Fourier transforming, we get

$$\left\langle n^{\alpha}(-\vec{q})n^{\beta}(\vec{q})\right\rangle_1 = \frac{\rho}{2}\delta^{\alpha\beta}. \tag{A.19}$$

We now calculate the magnetic-field correlator $\langle B(-\vec{q})B(\vec{q})\rangle$ using (A.16) along the q_z axis near the $\vec{G} = (002)$ pinch point. Since $\vec{G}$ is a reciprocal lattice point of the FCC lattice, both sublattices of the diamond lattice behave the same way as at $\vec{q} = 0$, but a relative phase $\vec{G}\cdot\vec{r}_\mu = \pi$ is introduced between them. In practice, this means that the $\langle n^{\pm}n^{\mp}\rangle$ correlators are the opposite of what they were near $\vec{q} = 0$, while the sign of $\langle n^{\pm}n^{\pm}\rangle$ remains unaffected. Equations (A.16, A.18, A.19) now yield

$$\langle b(-\vec{q})b(\vec{q})\rangle_{\mathrm{free}} = \frac{(2\pi)^2\rho}{4\sin^2(k/8)}\frac{k^2}{k^2+\xi_{\mathrm{D}}^{-2}} \approx \frac{64\pi^2}{K}\frac{1}{1+\xi_{\mathrm{D}}^2k^2}, \tag{A.20}$$

where $\vec{q} = (0, 0, 4\pi/a_0 + k)$. That is, the Debye plasma of visons introduces a Lorentzian blurring of the pinch points, with width $\xi_{\mathrm{D}}^{-1} \propto \sqrt{\rho}$.

A.3.2 Tightly bound dipoles

Since the dominant vison species at low temperatures is not the isolated vison, but a nearest-neighbour dipole of them, we need to derive $\langle b(-\vec{q})b(\vec{q})\rangle$ for a gas of such dipoles. Furthermore, the energy of a second-neighbour dipole is significantly smaller (by about $0.48g$) than predicted by the simple Coulomb approximation [Fig. 2.8(b)]. The Debye plasma approximation significantly underestimates the population of these dipoles, which must be corrected for explicitly. As the only effect of these associated dipoles is a correction to the polarisability of the emergent magnetic field, their interactions with one another can be neglected, so the only contribution to $\langle nn\rangle$ correlators comes from visons within the same dipole.

For nearest-neighbour ($\vec{r}_1 - \vec{r}_2 = a_0\langle 111\rangle/4$) dipoles, the real-space vison correlators

are thus

$$\langle n(\vec{r})n(\vec{r}')\rangle \propto \begin{cases} 2\rho & \vec{r}=\vec{r}' \\ -\rho/4 & \vec{r}'-\vec{r}=a_0\langle 111\rangle/4 \\ 0 & \text{otherwise,} \end{cases} \tag{A.21}$$

where ρ is the density of dipoles; the factor of $1/4$ is due to the four possible orientations of the dipole. Clearly, the two visons are on the same sublattice in the first line in (A.21) and on different ones in the second; therefore, the reciprocal-space correlators are

$$\left\langle n^{\alpha}(-\vec{q})n^{\beta}(\vec{q})\right\rangle = \rho \begin{pmatrix} 1 & -\gamma^* \\ -\gamma & 1 \end{pmatrix}^{\alpha\beta}, \tag{A.22}$$

from which the magnetic field correlator follows as

$$\langle b(-\vec{q})b(\vec{q})\rangle_{\mathrm{nn}} = \frac{\rho\pi^2}{4}\frac{2|\kappa|^2-\gamma^*\kappa^2-\gamma\kappa^{*2}}{1-|\gamma|^2}, \tag{A.23}$$

which is constant along the q_z axis. That is, the contribution of nearest-neighbour dipoles to the magnetic-field correlator is a uniform background that gradually submerges the pinch point.

For second-neighbour dipoles ($\vec{r}_1-\vec{r}_2=a_0\langle 110\rangle/2$), real space vison correlators are

$$\langle n(\vec{r})n(\vec{r}')\rangle \propto \begin{cases} 2\rho & \vec{r}=\vec{r}' \\ -\rho/6 & \vec{r}'-\vec{r}=a_0\langle 110\rangle/2 \\ 0 & \text{otherwise.} \end{cases} \tag{A.24}$$

Since the two visons of the dipole are on the same sublattice, the reciprocal space correlators are

$$\left\langle n^{\alpha}(-\vec{q})n^{\beta}(\vec{q})\right\rangle = \rho\left(1-\overbrace{\frac{1}{3}\sum_{i<j}\cos\frac{q_ia_0}{2}\cos\frac{q_ja_0}{2}}^{\Phi}\right)\delta^{\alpha\beta} \tag{A.25}$$

$$\langle b(-\vec{q})b(\vec{q})\rangle_{2\mathrm{nn}} = 8\rho\pi^2\frac{(1-\Phi)|\gamma^*\kappa-\kappa^*|^2}{\left(1-|\gamma|^2\right)^2}, \tag{A.26}$$

which is proportional to $\sin^2(q_za_0/8)$ along the q_z axis. This explains the small but significant cosine modulation of the lifting of the pinch point.

Besides the q_z axis, we considered the behaviour of the $\langle bb\rangle$ correlators around the pinch points, especially the one at $\vec{G}=(002)$. Vison-number correlators are analytic in all cases; however, the "response function" $\tilde{v}$ gives rise to pinch points perpendicular to

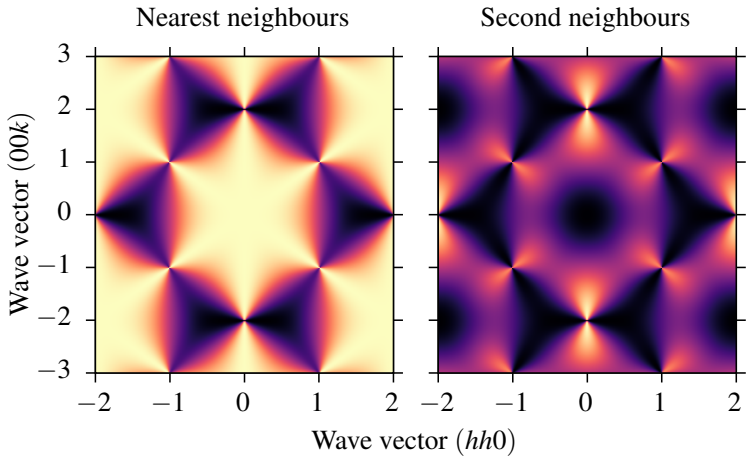

Figure A.1. Static correlation function $\langle b(-\vec{q})b(\vec{q})\rangle$ due to uniformly distributed bound pairs of visons at nearest-neighbour (left) and second-neighbour (right) distance. In the first case, the pinch-point pattern is complementary to that due to photons [Fig. 2.9(b)]. For second neighbours, the correlator is largest near the pinch points. This behaviour is reminiscent of the Debye plasma of dissociated visons, which gives rise to Lorentzian correlator peaks in the immediate vicinity of pinch points. Figure taken from Ref. 2. Perceptionally uniform colour map chosen following Ref. 125.

those due to photons (Fig. 2.9). The Debye-plasma contribution is peaked at the pinch point; its maximum matches the intensity of the photon contribution there, which results in Lorentzian pinch-point blurring [12, 76]. On the other hand, closely associated dipoles give rise to contributions proportional to the density throughout reciprocal space, which reduce the overall contrast of the pinch points. For completeness, the $\langle bb\rangle$ correlators produced by first- and second-neighbour dipoles, Eqs. (A.23) and (A.26), are plotted on the (hhk) plane in Fig. A.1.

A.4 Estimating the vison gap from thermal statistics

We discuss two methods to estimate the chemical potential of visons in a thermal ensemble, and derive expressions for both in terms of statistics of the vison number and the energy at a single temperature point, in a similar fashion to the fluctuation–dissipation theorem. We treat our model as a classical thermodynamic system with microstates of well-defined energy E_α and vison number N_α, weighted according to the partition function

$$Z = \sum_\alpha e^{-\beta E_\alpha + \zeta N_\alpha}, \tag{A.27}$$

where ζ is a fictitious chemical potential introduced to keep track of vison number; in the physical partition function, $\zeta = 0$.

The first approach estimates the energy cost of a single vison directly, which can formally be written as $\mu_E = \mathrm{d}E(N)/\mathrm{d}N$, where $E(N)$ is the mean energy of the system constrained to N visons. In the thermodynamic limit, this derivative is equivalent to the ratio of variations in $\langle E\rangle$ and $\langle N\rangle$ due to a changing chemical potential:

$$\mu_E \xrightarrow{N\to\infty} \frac{\partial E(\zeta)/\partial\zeta}{\partial N(\zeta)/\partial\zeta}. \tag{A.28}$$

Each of the two derivatives in (A.28) can be expressed in terms of derivatives of Z, which can in turn be rewritten in terms of statistics of E and N:

$$\left.\frac{\partial E}{\partial\zeta}\right|_\beta = -\frac{\partial^2 \log Z}{\partial\beta\partial\zeta} = \frac{1}{Z^2}\frac{\partial Z}{\partial\beta}\frac{\partial Z}{\partial\zeta} - \frac{1}{Z}\frac{\partial^2 Z}{\partial\beta\partial\zeta} = \langle NE\rangle - \langle N\rangle\langle E\rangle = \mathrm{cov}(N,E); \tag{A.29}$$

$$\left.\frac{\partial N}{\partial\zeta}\right|_\beta = \left.\frac{\partial^2 \log Z}{\partial\zeta^2}\right|_\beta = \frac{1}{Z}\left.\frac{\partial^2 Z}{\partial\zeta^2}\right|_\beta - \left(\frac{1}{Z}\left.\frac{\partial Z}{\partial\zeta}\right|_\beta\right)^2 = \langle N^2\rangle - \langle N\rangle^2 = \mathrm{var}\,N, \tag{A.30}$$

and therefore

$$\mu_E = \frac{\mathrm{cov}(E,N)}{\mathrm{var}\,N}. \tag{A.31}$$

Another estimate of the excitation gap is the local slope of the Arrhenius plot $\log N$ vs. $1/T$, at least at low temperatures where N has not saturated. However, this is not necessarily the gap of a single vison, but of whatever (possibly multi-vison) excitations are created thermally in the system. This slope is given by

$$\mu_{\text{Arrh.}} = -\frac{\mathrm{d}\log N}{\mathrm{d}\beta} = -\frac{1}{N}\frac{\mathrm{d}N}{\mathrm{d}\beta} = -\frac{1}{N}\frac{\partial^2 \log Z}{\partial\beta\partial\zeta} = \frac{\mathrm{cov}(E,N)}{N}.$$

Finally, it is instructive to consider $\mu_{\text{Arrh.}}/\mu_E = \mathrm{var}\,N/N$. As discussed before, if the dominant thermal (collective) excitation of the system consists of m visons, $\mu_{\text{Arrh.}}$ is expected to be $m\mu_E$ at low temperatures, and so $\mathrm{var}\,N/N \approx m$. We can obtain this last result directly by considering that the number $\tilde{N}$ of collective excitations obeys a Poisson distribution at low temperatures and so $\mathrm{var}\,\tilde{N} = \tilde{N}$. The result then follows from $N = m\tilde{N}$.

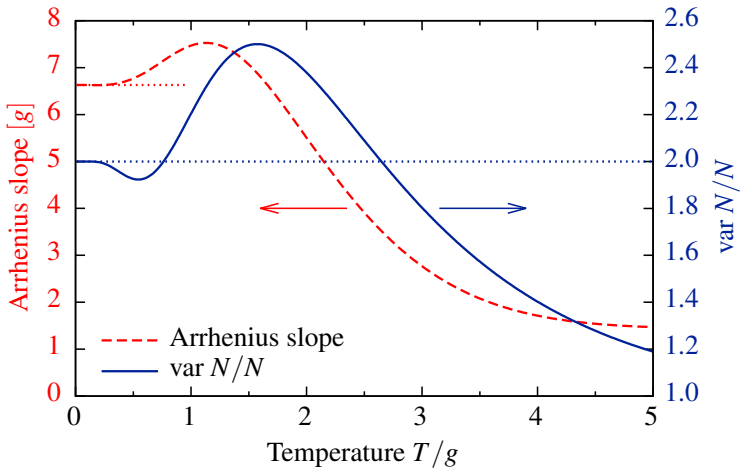

Figure A.2. Effective Arrhenius gap $\mu_{\text{Arrh.}} = -\mathrm{d}\log N/\mathrm{d}\beta$ (red dashed line) and var N/N (blue solid line) of visons in the semiquantitative model discussed in Appendix A.5. The zero-temperature energy cost of bare visons and nearest-neighbour vison pairs, E_1 and E_2, are given by the low-temperature simulations in §2.4.2. The model yields the expected zero-temperature limits of both quantities, as well as a qualitative temperature dependence similar to that seen in Fig. 2.12, for a wide range of the phenomenological parameters $E_{\text{int}}(T = 0)$ and m. They were chosen by hand to be $5g$ and 20, respectively, in order to achieve a good numerical agreement. Figure taken from Ref. 2.

A.5 Semiquantitative model of the partition function

The behaviour of vison number found in the simulations can be explained by the following three observations:

(i) The energy cost of a nearest-neighbour dipole of visons is smaller than that of a single isolated vison.

(ii) An isolated vison polarises its surroundings, reducing the energy cost of nearby, aligned dipoles even further.

(iii) The energy cost of visons reduces as temperature increases because highly excited photon modes "wash them out."

To demonstrate this, we construct a simplistic model of the partition function (A.27), which is semiquantitative at low temperatures and captures the salient features at large T. If we ignore interactions beyond nearest neighbours, the partition function factorises by diamond lattice sites: $Z = Z_1^{\mathcal{V}}$, where $\mathcal{V}$ is the number of such sites. Now, each site can host an isolated vison of either charge, or one half of a nearest-neighbour dipole. Assuming their energy cost is E_1 and E_2, respectively, we can write down a first approx-

imation to Z_1 as

$$Z_1 \approx 1 + 2e^{-\beta E_1 + \zeta} + 4e^{-\beta E_2 + 2\zeta}, \tag{A.32}$$

where $E_2 < E_1$ in line with observation (i).

Observation (ii) is implemented by changing the energy cost E_2 of a dipole near an isolated vison to $E_2 \pm E_{\text{int}}$ depending on its orientation. This changes the partition function of the link by

$$\Delta Z_1 \approx e^{2\zeta}\left(e^{-\beta(E_2+E_{\text{int}})} + e^{-\beta(E_2-E_{\text{int}})} - 2e^{-\beta E_2}\right) = 2e^{-\beta E_2 + 2\zeta}\left[\cosh(\beta E_{\text{int}}) - 1\right]. \tag{A.33}$$

In a crude approximation, we assume that each isolated vison introduces a fixed E_{int} to m nearby links and that the resulting ΔZ_1 can be factored into the partition function of the isolated vison. This gives

$$Z_1 \approx 1 + 4e^{-\beta E_2 + 2\zeta} + 2e^{-\beta E_1 + \zeta}(1 + \Delta Z_1)^m. \tag{A.34}$$

Observation (iii) is concerned with the strong interactions between a highly excited photon bath and the visons, a full treatment of which is a tall order. To estimate its effect on the vison thermodynamics, we propose a "Hartree–Fock approximation," where the effect of visons on the photon cloud is neglected and we assume that the quadratic theory that governs the photon modes at low temperatures [17] holds at arbitrary T. In this approximation, the energy associated with the gradient component b_{grad} of the magnetic field is

$$E(b_{\text{grad}}, T) = -g\langle\cos(b_{\text{grad}} + b_{\text{curl}})\rangle = -g\cos(b_{\text{grad}})e^{-T/(4g)}, \tag{A.35}$$

since b_{curl} on each plaquette has a Gaussian distribution of variance $T/(2g)$. That is, all energy scales associated with b_{grad}, and hence with the visons, are exponentially suppressed at high temperatures.

The partition function (A.33–A.35) can now be used to derive thermodynamic quantities: $\operatorname{var} N/N$ and the slope of the Arrhenius curve, $\mathrm{d}(\log N)/\mathrm{d}\beta$, are plotted in Fig. A.2. The results are in good qualitative agreement with the numerical simulation in Fig. 2.12. Saturation occurs at a higher temperature than in the full large-S treatment, albeit well below the zero-temperature energy cost of visons. More accurate estimates would likely follow from taking the emergent electric field and photon–photon interactions into account, but this is beyond the scope of this work.

B

Effective field theories for constrained quantum systems

B.1 Exact ground state at the large-S RK point

The ground state of the large-S RK Hamiltonian (3.9) at the RK point $V = J$ can be derived similarly to the $S = 1$ case (§3.2). We first write (3.9) in a similar form to (3.3):

$$H = J \sum_p \left(b_1^\dagger b_3^\dagger - b_2^\dagger b_4^\dagger\right)\left(b_1 b_3 - b_2 b_4\right) \tag{B.1}$$

on plaquettes of four links. Now, each term of the expectation value $\langle\psi|H|\psi\rangle$ is the norm of a state $(b_1 b_3 - b_2 b_4)|\psi\rangle$; therefore, H is positive semidefinite. Furthermore, if $(b_1 b_3 - b_2 b_4)|\psi\rangle$ vanishes for all plaquettes, $|\psi\rangle$ has zero energy, so it is the ground state of H.

To construct such a state, consider products of bosonic coherent states (3.10) on each link, projected onto the subspace that obeys the large-S constraint $\sum_{\ell\in+} b_\ell^\dagger b_\ell = S$ for all vertices:

$$|\psi\rangle = \hat{\mathcal{P}} \prod_\ell |\beta_\ell\rangle_\ell, \tag{B.2}$$

where $\hat{\mathcal{P}}$ is a projection operator onto the constrained subspace. We also define the operator $\hat{\mathcal{P}}'$ that projects onto the subspace where all vertices have S dimers, except for those around a particular plaquette, which have $(S-1)$. It follows that $b_1 b_3 \hat{\mathcal{P}} = \hat{\mathcal{P}}' b_1 b_3$ and $b_2 b_4 \hat{\mathcal{P}} = \hat{\mathcal{P}}' b_2 b_4$, as both pairs of annihilation operators remove one dimer from precisely these vertices. Therefore,

$$\begin{gathered} b_1 b_3|\psi\rangle = \hat{\mathcal{P}}' b_1 b_3 \prod_\ell |\beta_\ell\rangle_\ell = \hat{\mathcal{P}}' \beta_1\beta_3 \prod_\ell |\beta_\ell\rangle_\ell; \quad b_2 b_4|\psi\rangle = \hat{\mathcal{P}}' \beta_2\beta_4 \prod_\ell |\beta_\ell\rangle_\ell \\ (b_1 b_3 - b_2 b_4)|\psi\rangle = (\beta_1\beta_3 - \beta_2\beta_4)\hat{\mathcal{P}}' \prod_\ell |\beta_\ell\rangle_\ell. \end{gathered} \tag{B.3}$$

Therefore, if $\beta_1\beta_3 = \beta_2\beta_4$ for each plaquette, $|\psi\rangle$ is the exact ground state of the RK Hamiltonian (B.1) for arbitrary S. A trivial choice is setting all β to be equal, which concludes the construction.

In full generality, the constraint $\beta_1\beta_3 = \beta_2\beta_4$ around all plaquettes is satisfied if $\beta_\ell = \alpha_{v_1}\alpha_{v_2}$, where v_1 and v_2 are the endpoints of link ℓ, and the $\{\alpha_v\}$ are arbitrary complex numbers for each vertex. Now, using (B.2, 3.10), the ground state can be written in the number basis as

$$\begin{aligned}\langle\{n_\ell\}|\psi\rangle &= \prod_\ell e^{-|\beta_\ell|^2/2}\frac{\beta_\ell^{n_\ell}}{\sqrt{n_\ell!}} = \prod_\ell e^{-|\beta_\ell|^2/2}\frac{\alpha_1^{n_\ell}\alpha_2^{n_\ell}}{\sqrt{n_\ell!}} = \prod_\ell \frac{1}{\sqrt{n_\ell!}}\prod_v \alpha_v^S \prod_\ell e^{-|\beta_\ell|^2/2};\\ &\propto \prod_\ell \frac{1}{\sqrt{n_\ell!}},\end{aligned} \tag{B.4}$$

since n_ℓ sums to S around all vertices in every relevant configuration. That is, all valid choices of β in (B.2) lead to the same ground state up to irrelevant prefactors.

We finally note that in the $S = 1$ case, $n_\ell = 0, 1$ in all relevant configurations, so $\prod_\ell 1/\sqrt{n_\ell!}$ is equal in all of them: that is, Eq. (B.4) recovers the standard RK wave function (3.4) in this limit.

B.2 Instanton measure on the square lattice

Instantons appear naturally in the compact gauge theory as stationary trajectories of the action with ϕ changing by 2π on a given site between $\tau = \pm\infty$. Unlike the point-like instantons described in §3.4.3, these objects are smooth as a function of time, that is, they have a nontrivial instanton core. We expand the action to quadratic order around such solutions, similarly to the case $h = \phi \equiv 0$ shown in §3.4.2. The resulting fluctuation determinant appears in the probability of instantons as a pre-exponential factor.

First, we have to construct such a stationary instanton solution. Using (3.28, 3.33) with a single instanton event at $\vec{r} = \tau = 0$, we get the following quadratic action in terms of Fourier components:

$$\begin{aligned}\mathcal{S} = \int (\mathrm{d}\omega)(\mathrm{d}^2k)\Bigg\{&\omega h(\vec{k},\omega)\phi(-\vec{k},-\omega) + \frac{\mathcal{D}_0(\vec{k})}{2}h(\vec{k},\omega)h(-\vec{k},-\omega)\\ &+\frac{M}{2}\left[\phi(\vec{k},\omega) - \frac{2\pi i}{\omega}\right]\times \mathrm{c.c}\Bigg\},\end{aligned} \tag{B.5}$$

where we introduce $M = JS^2/8$ for brevity; i/ω is the Fourier transform of the Heaviside step function in τ. Since the different Fourier components are decoupled in this action,

we can minimise with respect to them separately, resulting in the stationary action

$$h_0(\vec{k},\omega) = \frac{2\pi i q M}{\omega^2 + M\mathcal{D}_0(\vec{k})} \qquad \phi_0(\vec{k},\omega) = \frac{2\pi i q}{\omega}\frac{M\mathcal{D}_0(\vec{k})}{\omega^2 + M\mathcal{D}_0(\vec{k})}. \tag{B.6}$$

In real space, this corresponds to a point-like instanton described in §3.4.3 together with a power-law decaying "instanton core". We should also note that (B.6) is not a stationary trajectory under the original action, but it becomes one if the $M(1-\cos\phi)$ potential term is replaced with the continued parabolic potential $V(\phi) = \frac{M}{2}\min_n(\phi - 2\pi n)^2$. Indeed, the only point where ϕ_0 reaches π is $\vec{r} = \tau = 0$, where the cusp in the potential is recovered by the Villain substitution (3.33). Therefore, we use $V(\phi)$ in what follows.

The action can now be expanded to quadratic order in δh and $\delta\phi$ around both the trivial trajectory $h = \phi \equiv 0$ and the instanton trajectory (B.6). δh can easily be integrated out in both cases, giving the following actions in $\delta\phi$:

$$\delta\mathcal{S}_0 = \frac{1}{2}\int (\mathrm{d}\omega)(\mathrm{d}^2k)\left(\frac{\omega^2}{\mathcal{D}_0(\vec{k})} + M\right)\delta\phi(\vec{k},\omega)\delta\phi(-\vec{k},-\omega) \tag{B.7}$$

$$\delta\mathcal{S}_\mathrm{i} = \delta\mathcal{S}_0 - \frac{\pi M}{\partial_\tau\phi_0(\vec{r}=\tau=0)}[\delta\phi(\vec{r}=\tau=0)]^2, \tag{B.8}$$

where the additional term in (B.8) corresponds to the cusp of the continued parabolic potential at $\phi = \pi$ reached at $\vec{r} = \tau = 0$. The most important difference between the two actions is that $\delta\mathcal{S}_\mathrm{i}$ has a zero mode $\psi = \partial_\tau\phi_0$, which corresponds to the continuous time-translation symmetry of the setup. For such modes, the usual contribution to the partition function, $e^{-\mathcal{S}_\mathrm{cl}}(\det K)^{-1/2}$, is replaced by

$$\mathrm{d}\tau\sqrt{\frac{\langle\psi|\psi\rangle}{2\pi}}e^{-\mathcal{S}_\mathrm{cl}}(\det\tilde{K})^{-1/2}, \tag{B.9}$$

where $\mathcal{S}_\mathrm{cl}$ is the action due to the stationary instanton (3.35), and $\tilde{K}$ is the fluctuation kernel of (B.8) restricted to non-zero modes. Suppose $|\psi\rangle$ is proportional to a basis vector (this can always be achieved using a unitary transformation on the kernel K). Then, $\tilde{K}$ is the principal minor of K that excludes the row and column of ψ. $\det\tilde{K}$ is thus a cofactor of the full kernel, and we formally have

$$\det\tilde{K} = \langle\tilde{\psi}|K^{-1}|\tilde{\psi}\rangle\det K \tag{B.10}$$

by the cofactor formula for matrix inversion, where $|\tilde{\psi}\rangle$ is the normalised zero mode $|\psi\rangle/\sqrt{\langle\psi|\psi\rangle}$.

In a matrix language, K is the sum of the kernel K_0 of (B.7) and a dyad $-\lambda|v\rangle\langle v|$, where $|v\rangle$ corresponds to the δ-function at $\vec{r} = \tau = 0$ implied in (B.8) in an arbitrary basis. For such a matrix, we have the following:

(i) $\det K = \det K_0(1 - \lambda\langle v|K_0^{-1}|v\rangle)$;

(ii) If $1 - \lambda\langle v|K_0^{-1}|v\rangle = 0$, $K_0^{-1}|v\rangle$ is an eigenvector of K with zero eigenvalue;

(iii) $K^{-1} = K_0^{-1} + \dfrac{\lambda K_0^{-1}|v\rangle\langle v|K_0^{-1}}{1 - \lambda\langle v|K_0^{-1}|v\rangle}$.

The first statement can be proved by inserting factors of $K_0^{1/2}K_0^{-1/2}$ into the definition of K; the other two are straightforward to verify. Now, $\det\tilde{K}$ follows as

$$\det\tilde{K} = \left\langle\tilde{\psi}_0\middle|\left(K_0^{-1} + \frac{\lambda K_0^{-1}|v\rangle\langle v|K_0^{-1}}{1 - \lambda\langle v|K_0^{-1}|v\rangle}\right)\middle|\tilde{\psi}_0\right\rangle \times (1 - \lambda\langle v|K_0^{-1}|v\rangle) \times \det K_0 \tag{B.11}$$

$$= \langle\tilde{\psi}_0|\lambda K_0^{-1}|v\rangle\langle v|K_0^{-1}|\tilde{\psi}_0\rangle \det K_0$$

$$= \lambda\langle v|K_0^{-2}|v\rangle \det K_0. \tag{B.12}$$

In the first line, we substitute statements (i) and (iii) into (B.10); in the second, we note that $1 - \lambda\langle v|K_0^{-1}|v\rangle = 0$ in our case, so only the second term of K^{-1} gives any contribution.[†] Finally, we use that $|\tilde{\psi}_0\rangle$ is the normalised zero mode, so, by statement (ii), it is $K_0^{-1}|v\rangle/\sqrt{\langle v|K_0^{-2}|v\rangle}$. Altogether, the measure of the instanton solutions relative to that of the instanton-free solution, $(\det K_0)^{-1/2}$, is

$$\mathrm{d}\tau\sqrt{\frac{\langle\psi|\psi\rangle}{2\pi\lambda\langle v|K_0^{-2}|v\rangle}}e^{-\mathcal{S}_{\mathrm{cl}}}. \tag{B.13}$$

In the $(\vec{k},\omega)$ basis, K_0 is positive definite and diagonal, $K_0(\vec{k},\omega) = \omega^2/\mathcal{D}_0(\vec{k}) + M$, $\lambda = 2\pi M/[\partial_\tau\phi_0(0,0)]$, and $v(\vec{k},\omega) = 1$ (the Fourier transform of a δ-function at the origin). It is easy to verify that $1 - \lambda\langle v|K_0^{-1}|v\rangle = 0$, that is, K indeed has a zero mode; furthermore, $K_0^{-1}|v\rangle \propto \partial_\tau\phi_0$ as expected. Now, substituting into (B.13) gives that the measure of instanton solutions of a given sign is $\mu I\mathrm{d}\tau$, where $I = e^{-\mathcal{S}_0}$ and

$$\mu = \sqrt{M\pi\int(\mathrm{d}^2k)\omega(\vec{k})}; \tag{B.14}$$

at the RK point on the square lattice, $\mu = JS^{3/2}\sqrt{\pi/8}$.

[†]Note that, since $1 - \lambda\langle v|K_0^{-1}|v\rangle = 0$ due to the zero mode, both terms in (B.11) are ill-defined. However, both are finite for any other λ, so we can formally evaluate it elsewhere and take the limit of λ going to its physical value. The divergence due to $1 - \lambda\langle v|K_0^{-1}|v\rangle = 0$ cancels, so (B.12) is indeed the correct limit.

C

Neural network wave functions and the sign problem

C.1 Stochastic reconfiguration

Stochastic reconfiguration proceeds by approximating the imaginary-time evolution of the trial wave function using Monte-Carlo sampling. Namely, given $|\Psi_\theta\rangle$, we want to find a new set of the real parameters[†] $\theta' = \theta + \delta\theta$ such that $|\Psi'\rangle = |\Psi_{\theta'}\rangle$ is a good approximation to

$$|\Psi'_{\text{exact}}\rangle = e^{-\eta H}|\Psi_\theta\rangle \approx (1 - \eta H)|\Psi_\theta\rangle, \tag{C.1}$$

where η is a small positive number that plays the role of the learning rate in machine learning language. Since we only want to project out all excited states, the error introduced by expanding $e^{-\eta H}$ in (C.1) is irrelevant. $|\Psi'\rangle$ is optimised by maximising the overlap of the (unnormalised) wave functions $|\Psi'_{\text{exact}}\rangle$ and $|\Psi'\rangle$,

$$|C|^2 = \frac{\langle\Psi'_{\text{exact}}|\Psi'\rangle\langle\Psi'|\Psi'_{\text{exact}}\rangle}{\langle\Psi'_{\text{exact}}|\Psi'_{\text{exact}}\rangle\langle\Psi'|\Psi'\rangle}. \tag{C.2}$$

To linear order in both η and $\delta\theta$, the condition $\partial_{\theta'_k}|C|^2 = 0$ leads to

$$\begin{aligned}\sum_j \delta\theta_j \,\mathrm{Re}\left[\frac{\langle\partial_{\theta_j}\Psi|\partial_{\theta_k}\Psi\rangle}{\langle\Psi|\Psi\rangle} - \frac{\langle\partial_{\theta_j}\Psi|\Psi\rangle}{\langle\Psi|\Psi\rangle}\frac{\langle\Psi|\partial_{\theta_k}\Psi\rangle}{\langle\Psi|\Psi\rangle}\right] \\ = \eta\,\mathrm{Re}\left[\frac{\langle\Psi|H|\partial_{\theta_k}\Psi\rangle}{\langle\Psi|\Psi\rangle} - \frac{\langle\Psi|H|\Psi\rangle}{\langle\Psi|\Psi\rangle}\frac{\langle\Psi|\partial_{\theta_k}\Psi\rangle}{\langle\Psi|\Psi\rangle}\right].\end{aligned} \tag{C.3}$$

[†]Equivalent expressions can be derived for trial wave functions that are (piecewise) analytic functions of complex parameters [38, 203]. The result is identical to (C.7), omitting the real-part signs.

In order to find $\delta\theta$ numerically, we rewrite the expectation values in (C.3) as Monte-Carlo averages with respect to the quantum probability distribution $p(\boldsymbol{\sigma}) = \left|\langle\boldsymbol{\sigma}|\Psi\rangle\right|^2/\langle\Psi|\Psi\rangle$. This can readily be done by inserting a resolution of the identity; for example, we have

$$\begin{aligned}\frac{\langle\Psi|H|\partial_{\theta_k}\Psi\rangle}{\langle\Psi|\Psi\rangle} &= \sum_{\boldsymbol{\sigma}} \frac{\langle\Psi|H|\boldsymbol{\sigma}\rangle\,\langle\boldsymbol{\sigma}|\partial_{\theta_k}\Psi\rangle}{\langle\Psi|\Psi\rangle} = \sum_{\boldsymbol{\sigma}} \frac{\left|\langle\boldsymbol{\sigma}|\Psi\rangle\right|^2}{\langle\Psi|\Psi\rangle}\frac{\langle\Psi|H|\boldsymbol{\sigma}\rangle}{\langle\Psi|\boldsymbol{\sigma}\rangle}\frac{\langle\boldsymbol{\sigma}|\partial_{\theta_k}\Psi\rangle}{\langle\boldsymbol{\sigma}|\Psi\rangle} \\ &= \sum_{\boldsymbol{\sigma}} p(\boldsymbol{\sigma}) E^*_{\mathrm{loc}}(\boldsymbol{\sigma}) O_k(\boldsymbol{\sigma}),\end{aligned} \tag{C.4}$$

where we introduce

$$O_k(\boldsymbol{\sigma}) = \frac{\langle\boldsymbol{\sigma}|\partial_{\theta_k}\Psi\rangle}{\langle\boldsymbol{\sigma}|\Psi\rangle} = \partial_{\theta_k}\log\langle\boldsymbol{\sigma}|\Psi\rangle \tag{C.5}$$

and the local energy

$$E_{\mathrm{loc}}(\boldsymbol{\sigma}) = \frac{\langle\boldsymbol{\sigma}|H|\Psi\rangle}{\langle\boldsymbol{\sigma}|\Psi\rangle}. \tag{C.6}$$

The expectation value (C.4) can now be estimated as the Monte-Carlo average of $E^*_{\mathrm{loc}} O_k$ for samples distributed according to $p(\boldsymbol{\sigma})$; the others follow analogously, resulting in

$$\sum_j \underbrace{\operatorname{Re}\operatorname{cov}(O_j, O_k)}_{S_{kj}}\,\delta\theta_j = -\eta\,\underbrace{\operatorname{Re}\operatorname{cov}(E_{\mathrm{loc}}, O_k)}_{\partial_{\theta_k}E(\theta)}, \tag{C.7}$$

as stated in (4.3). Solving this equation for $\delta\theta_k$ requires inverting the covariance matrix S, which, while positive semidefinite, tends to be ill-conditioned even for a large number of Monte-Carlo samples [38]. This can be resolved either by using the pseudoinverse, or, more commonly, by adding a small positive constant to the diagonal entries in order to make the matrix invertible.

C.2 Estimating observables

Once the wave function had converged, the estimates of cumulative distribution functions plotted in Figs. 4.5 and 4.8 were generated by drawing 100 000 samples out of the probability distribution $p(\boldsymbol{\sigma}) = \left|\langle\boldsymbol{\sigma}|\Psi\rangle\right|^2/\langle\Psi|\Psi\rangle$ and sorting their phases.

Analogous to the estimate (C.6) of variational energy, the expectation value of any operator A can be evaluated as the Monte-Carlo average of the local estimates

$$A_{\mathrm{loc}}(\boldsymbol{\sigma}) = \frac{\langle\boldsymbol{\sigma}|A|\Psi\rangle}{\langle\boldsymbol{\sigma}|\Psi\rangle} \tag{C.8}$$

with respect to the quantum probability distribution $p(\boldsymbol{\sigma}) = \left|\langle\boldsymbol{\sigma}|\Psi\rangle\right|^2/\langle\Psi|\Psi\rangle$. NetKet

provides a facility for evaluating such expectation values within the SR protocol: this was used to estimate the variational energy, $\langle \vec{S}^2 \rangle$, and the spin correlators (4.20) using 1 000 000 Monte Carlo samples. We exploited the translation invariance of the wave function to rewrite the latter two as

$$S^2(\vec{q}) = \frac{1}{N+2} \sum_i \langle \vec{\sigma}_i \cdot \vec{\sigma}_0 \rangle e^{i\vec{q}\cdot\vec{r}_i}; \tag{C.9}$$

$$\langle \vec{S}^2 \rangle = N \sum_i \langle \vec{\sigma}_i \cdot \vec{\sigma}_0 \rangle, \tag{C.10}$$

where both sums include $i = 0$ (note, however, that $\vec{\sigma}_0 \cdot \vec{\sigma}_0 \equiv 3/4$).

We checked furthermore whether the converged wave functions obey parity and point-group symmetries. The parity operator $\mathcal{P} = \prod_i \sigma_i^x$ commutes with all point-group symmetries as well as the Hamiltonian: therefore, the parity of a wave function is fully characterised by the expectation value $\langle \mathcal{P} \rangle$, without any possibility of symmetry-protected degeneracies. By contrast, the non-Abelian point group D_4 does give rise to such degeneracies, limiting the usefulness of plain symmetry-operator expectation values. Instead, we evaluated the statistical weight of eigenstates that transform according to the different irreps α of D_4 using the projection operators [378]

$$\hat{P}_\alpha = \frac{d_\alpha}{|D_4|} \sum_{g \in D_4} \chi_\alpha(g) \hat{g}, \tag{C.11}$$

where d_α and χ_α are the dimension and characters of the irrep, respectively, and $|D_4| = 8$. The weight of each irrep is given by

$$w_\alpha = \frac{\langle \psi | \hat{P}_\alpha^\dagger \hat{P}_\alpha | \psi \rangle}{\langle \psi | \psi \rangle} = \frac{\langle \psi | \hat{P}_\alpha | \psi \rangle}{\langle \psi | \psi \rangle} = \langle \hat{P}_\alpha \rangle,$$

where we used the fact that the projector (C.11) is Hermitian and squares to itself. Both w_α and $\langle \mathcal{P} \rangle$ were evaluated using 4 000 000 Monte Carlo samples in (C.8). [We used different samples here, as NetKet does not offer a simple implementation of these symmetry operators, so it proved more expedient to sample $p(\boldsymbol{\sigma})$ directly.]

Converged variational energies and all other observables are reported in Tables 4.1 and 4.2, respectively.

C.3 Average signs in exact diagonalisation

To estimate the average Marshall-adjusted sign of the true ground state of the 10×10 frustrated model, we obtained the exact ground states for 4×4 and 4×6 lattices using

	$\langle s \rangle$	$\langle s \rangle_{\mathrm{MSR}}$
4×4	3.53×10^{-2}	0.9745
4×6	3.87×10^{-3}	0.9650
10×10	$\approx 3 \times 10^{-12}$	≈ 0.88

Table C.1. Average sign $\langle s \rangle$ and Marshall-adjusted sign $\langle s \rangle_{\mathrm{MSR}}$ of the ground state of the $J_2/J_1 = 0.5$ square lattice HAFM for 4×4 and 4×6 lattices, calculated by exact diagonalisation. The 10×10 lattice is extrapolated from these, assuming that $\langle s \rangle$ decays exponentially in the number of spins [185].

the Lanczos algorithm as implemented in SciPy [379] and calculated the average sign

$$\langle s \rangle_\psi = \left| \sum_{\boldsymbol{\sigma}} p_\psi(\boldsymbol{\sigma}) \frac{\langle \boldsymbol{\sigma} | \psi \rangle}{\left| \langle \boldsymbol{\sigma} | \psi \rangle \right|} \right| = \left| \frac{\sum_{\boldsymbol{\sigma}} \left| \langle \boldsymbol{\sigma} | \psi \rangle \right| \langle \boldsymbol{\sigma} | \psi \rangle}{\sum_{\boldsymbol{\sigma}} \left| \langle \boldsymbol{\sigma} | \psi \rangle \right|^2} \right| \tag{C.12}$$

for both the original and the Marshall-adjusted ground states, $|\mathrm{GS}\rangle$ and $\prod_{i \in A} \sigma_i^z |\mathrm{GS}\rangle$. These average signs are given in Table C.1, together with an extrapolation to the 10×10 lattice, assuming an exponential decay of both $\langle s \rangle$ [185]. The average sign of the original wave function decays extremely fast and becomes negligibly small for our lattice size. By contrast, the average Marshall-adjusted sign remains close, but clearly distinct from, one. For a 10×10 lattice, the expected average sign is ≈ 0.88: since this is the difference of the statistical weights of positive- and negative-amplitude states, we expect those to be about 94% and 6%, respectively.

D

Monte-Carlo simulations of $Ho_2Ir_2O_7$

As described in §5.1.2, the behaviour of holmium spins in a single iridium AIAO domain of $Ho_2Ir_2O_7$ is captured by the dipolar spin-ice Hamiltonian [68] with an additional local $\langle 111 \rangle$ field to represent the coupling to the ordered iridium moments [69]:

$$\mathcal{H} = \frac{J}{3} \sum_{\langle ij \rangle} \sigma_i \sigma_j + D\ell^3 \sum_{ij} \sigma_i \sigma_j \left[\frac{\hat{e}_i \cdot \hat{e}_j}{r_{ij}^3} - \frac{3(\hat{e}_i \cdot \vec{r}_{ij})(\hat{e}_j \cdot \vec{r}_{ij})}{r_{ij}^5} \right] \pm h_{\text{loc}} \sum_i \sigma_i - \mu_{\text{Ho}} \mu_0 \vec{H} \cdot \sum_i \sigma_i \hat{e}_i, \tag{D.1}$$

where the last term explicitly includes the external magnetic field $\vec{H}$. Following experiments on $Ho_2Ti_2O_7$ [61], the magnitude μ_{Ho} of holmium moments was taken to be $10.0\,\mu_B$. Given the distance $\ell = a_0/\sqrt{8} = 3.5988\,\text{Å}$ between nearest-neighbour holmium atoms [69], this implies a dipolar interaction strength of $D = \mu_0 \mu_{\text{Ho}}^2/(4\pi\ell^3) \approx 1.23\,\text{K}$. The nearest-neighbour coupling was set to $J = -1.56\,\text{K}$ based on $Ho_2Ti_2O_7$ [61]; h_{loc} was set to $3.5\,\text{K}$ by trial and error to ensure a good correspondence between the experimental and simulated magnetisation curves in [100] fields.[†]

The effective model (D.1) was simulated on $6 \times 6 \times 6$ cubic unit cells of the pyrochlore lattice (3 456 holmium spins); h_{loc} was chosen positive throughout the sample to simulate the interior of an iridium AIAO domain.[‡] Periodic boundary conditions for the dipolar interaction were enforced using Ewald summation [127]; the same allows for introducing a demagnetising factor γ, which effectively changes the external magnetic field $\vec{H}$ to $\vec{H}_{\text{eff}} = \vec{H} - \gamma\vec{M}$, where $\vec{M}$ is the net magnetisation of the sample. This way, the simulations can be compared directly to the experimental data that are not corrected for demagnetisation effects. Due to the irregular octahedral shape of the sample, accounting for demagnetisation accurately is a tall order; as an approximation, we set $\gamma = 1/3$,

[†] No hysteresis was observed along this direction, suggesting that the response of type *A* and *B* domains are identical, so the experimental data can be used to fit our single-domain simulations.

[‡] In the hysteretic case, the response of the polycrystalline domain structure was approximated using weighted averages of the simulated pure type *A* and *B* domains, see §5.3.1.

consistent with a spherical sample [380]. We used single spin-flip simulations, which allowed us to study time evolution as well.[†]

[†]The magnetic field was changed at a rate of 0.2 Oe per MC sweep. In $Dy_2Ti_2O_7$, one sweep corresponds to about a millisecond in real time [272], so this sweep rate is equivalent to about 200 Oe/s, slightly faster than the fastest experimental rate in Fig. 5.7(a). In fact, no dynamical hysteresis was observed in the simulations up to about 0.1 T/s.

References

[1] A. Szabó, G. Goldstein, C. Castelnovo, and A. M. Tsvelik, Phys. Rev. B **100**, 085113 (2019).

[2] A. Szabó and C. Castelnovo, Phys. Rev. B **100**, 014417 (2019).

[3] A. Szabó and C. Castelnovo, Phys. Rev. Research **2**, 033075 (2020).

[4] M. J. Pearce, K. Götze, A. Szabó, T. S. Sikkenk, M. R. Lees, A. T. Boothroyd, D. Prabhakaran, C. Castelnovo, and P. A. Goddard, arXiv:2102.04483 (2021).

[5] A. Szabó and U. Schneider, Phys. Rev. B **98**, 134201 (2018).

[6] A. Szabó and U. Schneider, Phys. Rev. B **101**, 014205 (2020).

[7] L. Savary and L. Balents, Rep. Prog. Phys. **80**, 016502 (2017).

[8] G. H. Wannier, Phys. Rev. **79**, 357 (1950).

[9] R. J. Baxter, *Exactly Solved Models in Statistical Mechanics* (Academic Press, London, 1982).

[10] L. Pauling, J. Am. Chem. Soc. **57**, 2680 (1935).

[11] E. H. Lieb, Phys. Rev. **162**, 162 (1967).

[12] C. L. Henley, Phys. Rev. B **71**, 014424 (2005).

[13] C. L. Henley, Annu. Rev. Condens. Matter Phys. **1**, 179 (2010).

[14] D. A. Huse, W. Krauth, R. Moessner, and S. L. Sondhi, Phys. Rev. Lett. **91**, 167004 (2003).

[15] M. Hermele, M. P. A. Fisher, and L. Balents, Phys. Rev. B **69**, 064404 (2004).

[16] A. M. Polyakov, Phys. Lett. B **59**, 82 (1975).

[17] M. P. Kwasigroch, B. Douçot, and C. Castelnovo, Phys. Rev. B **95**, 134439 (2017).

[18] J. B. Kogut, Rev. Mod. Phys. **51**, 659 (1979).

[19] N. Shannon, O. Sikora, F. Pollmann, K. Penc, and P. Fulde, Phys. Rev. Lett. **108**, 067204 (2012).

[20] R. J. Baxter, Phys. Rev. Lett. **26**, 832 (1971).

[21] A. Y. Kitaev, Ann. Phys. **303**, 2 (2003).

[22] T. H. Hansson, V. Oganesyan, and S. L. Sondhi, Ann. Phys. **313**, 497 (2004).

[23] A. Hamma, R. Ionicioiu, and P. Zanardi, Phys. Rev. A **71**, 022315 (2005).

[24] A. Kitaev and J. Preskill, Phys. Rev. Lett. **96**, 110404 (2006).

[25] M. Levin and X.-G. Wen, Phys. Rev. Lett. **96**, 110405 (2006).

[26] R. Moessner and S. L. Sondhi, Prog. Theor. Phys. Suppl. **145**, 37 (2002).

[27] P. W. Anderson, Mater. Res. Bull. **8**, 153 (1973).

[28] P. Fazekas and P. W. Anderson, Philos. Mag. **30**, 423 (1974).

[29] P. W. Anderson, Science **235**, 1196 (1987).

[30] S. A. Kivelson, D. S. Rokhsar, and J. P. Sethna, Phys. Rev. B **35**, 8865(R) (1987).

[31] G. Baskaran, Z. Zou, and P. W. Anderson, Solid State Commun. **63**, 973 (1987).

[32] A. Auerbach, *Interacting Electrons and Quantum Magnetism* (Springer-Verlag, New York, 1994).

[33] I. Affleck and J. B. Marston, Phys. Rev. B **37**, 3774 (1988).

[34] I. Affleck, Z. Zou, T. Hsu, and P. W. Anderson, Phys. Rev. B **38**, 745 (1988).

[35] G. Baskaran and P. W. Anderson, Phys. Rev. B **37**, 580 (1988).

[36] X.-G. Wen, Phys. Rev. B **65**, 165113 (2002).

[37] X.-G. Wen, *Quantum Field Theory of Many-Body Systems: From the Origin of Sound to an Origin of Light and Electrons*, Oxford Graduate Texts (Oxford University Press, Oxford, 2007).

[38] F. Becca and S. Sorella, *Quantum Monte Carlo Approaches for Correlated Systems* (Cambridge University Press, Cambridge, 2017).

[39] M. C. Gutzwiller, Phys. Rev. Lett. **10**, 159 (1963).

[40] M. I. Aroyo, ed., *Space-group symmetry*, 2nd ed., International Tables for Crystallography, Vol. A (Wiley, New York, 2016).

[41] C. R. Wiebe and A. M. Hallas, APL Mater. **3**, 041519 (2015).

[42] F. Queyroux, C. R. Seances Acad. Sci., Ser. C **259**, 1527 (1964).

[43] K. Momma and F. Izumi, J. Appl. Crystallogr. **44**, 1272 (2011).

[44] J. Jensen and A. R. Mackintosh, *Rare Earth Magnetism* (Clarendon Press, Oxford, 1991).

[45] S. Rosenkranz, A. P. Ramirez, A. Hayashi, R. J. Cava, R. Siddharthan, and B. S. Shastry, J. Appl. Phys. **87**, 5914 (2000).

[46] Y. Jana, A. Sengupta, and D. Ghosh, J. Magn. Magn. Mater. **248**, 7 (2002).

[47] K. Kimura, S. Nakatsuji, J.-J. Wen, C. Broholm, M. B. Stone, E. Nishibori, and H. Sawa, Nat. Comm. **4**, 1934 (2013).

[48] H. A. Kramers, Proc. Amsterdam Acad. **33**, 959 (1930).

[49] E. Wigner, Nachr. Ges. Wiss. Göttingen **31**, 546 (1932).

[50] M. J. Gingras and P. A. McClarty, Rep. Prog. Phys. **77**, 056501 (2014).

[51] L. D. C. Jaubert, O. Benton, J. G. Rau, J. Oitmaa, R. R. P. Singh, N. Shannon, and M. J. Gingras, Phys. Rev. Lett. **115**, 267208 (2015).

[52] S. T. Bramwell and M. J. Gingras, Science **294**, 1495 (2001).

[53] C. Castelnovo, R. Moessner, and S. L. Sondhi, Annu. Rev. Condens. Matter Phys. **3**, 35 (2012).

[54] W. F. Giauque and M. F. Ashley, Phys. Rev. **43**, 81 (1933).

[55] W. F. Giauque and J. W. Stout, J. Am. Chem. Soc. **58**, 1144 (1936).

[56] J. D. Bernal and R. H. Fowler, J. Chem. Phys. **1**, 515 (1933).

[57] P. W. Anderson, Phys. Rev. **102**, 1008 (1956).

[58] M. J. Harris, S. T. Bramwell, D. F. McMorrow, T. Zeiske, and K. W. Godfrey, Phys. Rev. Lett. **79**, 2554 (1997).

[59] A. P. Ramirez, A. Hayashi, R. J. Cava, R. Siddharthan, and B. S. Shastry, Nature **399**, 333 (1999).

[60] S. V. Isakov, K. Gregor, R. Moessner, and S. L. Sondhi, Phys. Rev. Lett. **93**, 167204 (2004).

[61] S. T. Bramwell, M. J. Harris, B. C. den Hertog, M. J. Gingras, J. S. Gardner, D. F. McMorrow, A. R. Wildes, A. L. Cornelius, J. D. M. Champion, R. G. Melko, and T. Fennell, Phys. Rev. Lett. **87**, 047205 (2001).

[62] C. Castelnovo, R. Moessner, and S. L. Sondhi, Nature **451**, 42 (2008).

[63] C. Castelnovo, R. Moessner, and S. L. Sondhi, Phys. Rev. B **84**, 144435 (2011).

[64] Y. Levin, Rep. Prog. Phys. **65**, 1577 (2002).

[65] D. J. P. Morris, D. A. Tennant, S. A. Grigera, B. Klemke, C. Castelnovo, R. Moessner, C. Czternasty, M. Meissner, K. C. Rule, J.-U. Hoffmann, K. Kiefer, S. Gerischer, D. Slobinsky, and R. S. Perry, Science **326**, 411 (2009).

[66] T. Fennell, P. P. Deen, A. R. Wildes, K. Schmalzl, D. Prabhakaran, A. T. Boothroyd, R. J. Aldus, D. F. McMorrow, and S. T. Bramwell, Science **326**, 415 (2009).

[67] H. Kadowaki, N. Doi, Y. Aoki, Y. Tabata, T. J. Sato, J. W. Lynn, K. Matsuhira, and Z. Hiroi, J. Phys. Soc. Jpn. **78**, 103706 (2009).

[68] B. C. den Hertog and M. J. Gingras, Phys. Rev. Lett. **84**, 3430 (2000).

[69] E. Lefrançois, V. Cathelin, E. Lhotel, J. Robert, P. Lejay, C. V. Colin, B. Canals, F. Damay, J. Ollivier, B. Fåk, L. C. Chapon, R. Ballou, and V. Simonet, Nat. Comm. **8**, 209 (2017).

[70] R. G. Melko and M. J. Gingras, J. Phys. Condens. Matter **16**, R1277 (2004).

[71] K. Matsuhira, Y. Hinatsu, K. Tenya, and T. Sakakibara, J. Phys. Condens. Matter **12**, L649 (2000).

[72] K. Matsuhira, Y. Hinatsu, and T. Sakakibara, J. Phys. Condens. Matter **13**, L737 (2001).

[73] J. Snyder, B. G. Ueland, J. S. Slusky, H. Karunadasa, R. J. Cava, and P. Schiffer, Phys. Rev. B **69**, 064414 (2004).

[74] P. A. M. Dirac, Proc. R. Soc. London A **133**, 60 (1931).

[75] S. V. Isakov, R. Moessner, and S. L. Sondhi, Phys. Rev. Lett. **95**, 217201 (2005).

[76] A. Sen, R. Moessner, and S. L. Sondhi, Phys. Rev. Lett. **110**, 107202 (2013).

[77] H. Aoki, T. Sakakibara, K. Matsuhira, and Z. Hiroi, J. Phys. Soc. Jpn. **73**, 2851 (2004).

[78] S. T. Bramwell, S. R. Giblin, S. Calder, R. J. Aldus, D. Prabhakaran, and T. Fennell, Nature **461**, 956 (2009).

[79] F. K. Kirschner, F. Flicker, A. Yacoby, N. Y. Yao, and S. J. Blundell, Phys. Rev. B **97**, 140402 (2018).

[80] R. Dusad, F. K. K. Kirschner, J. C. Hoke, B. R. Roberts, A. Eyal, F. Flicker, G. M. Luke, S. J. Blundell, and J. C. S. Davis, Nature **571**, 234 (2019).

[81] C. A. Watson, I. Sochnikov, J. R. Kirtley, R. J. Cava, and K. A. Moler, arXiv:1903.11465 (2019).

[82] B. Cabrera, Phys. Rev. Lett. **48**, 1378 (1982).

[83] J. Knolle and R. Moessner, Annu. Rev. Condens. Matter Phys. **10**, 451 (2019).

[84] K. A. Ross, L. Savary, B. D. Gaulin, and L. Balents, Phys. Rev. X **1**, 021002 (2011).

[85] L. Savary and L. Balents, Phys. Rev. Lett. **108**, 037202 (2012).

[86] L. Savary and L. Balents, Phys. Rev. B **87**, 205130 (2013).

[87] S. Lee, S. Onoda, and L. Balents, Phys. Rev. B **86**, 104412 (2012).

[88] M. Taillefumier, O. Benton, H. Yan, L. D. C. Jaubert, and N. Shannon, Phys. Rev. X **7**, 041057 (2017).

[89] D. S. Rokhsar and S. A. Kivelson, Phys. Rev. Lett. **61**, 2376 (1988).

[90] A. Banerjee, S. V. Isakov, K. Damle, and Y. B. Kim, Phys. Rev. Lett. **100**, 047208 (2008).

[91] Y. Kato and S. Onoda, Phys. Rev. Lett. **115**, 077202 (2015).

[92] R. Moessner and J. T. Chalker, Phys. Rev. Lett. **80**, 2929 (1998).

[93] P. H. Conlon and J. T. Chalker, Phys. Rev. Lett. **102**, 237206 (2009).

[94] O. Benton, L. D. C. Jaubert, R. R. P. Singh, J. Oitmaa, and N. Shannon, Phys. Rev. Lett. **121**, 067201 (2018).

[95] O. Benton, O. Sikora, and N. Shannon, Phys. Rev. B **86**, 075154 (2012).

[96] S. Nakosai and S. Onoda, J. Phys. Soc. Jpn. **88**, 053701 (2019).

[97] M. P. Kwasigroch, Phys. Rev. B **102**, 125113 (2020).

[98] C.-J. Huang, Y. Deng, Y. Wan, and Z. Y. Meng, Phys. Rev. Lett. **120**, 167202 (2018).

[99] Y. Tokiwa, T. Yamashita, D. Terazawa, K. Kimura, Y. Kasahara, T. Onishi, Y. Kato, M. Halim, P. Gegenwart, T. Shibauchi, S. Nakatsuji, E.-G. Moon, and Y. Matsuda, J. Phys. Soc. Jpn. **87**, 064702 (2018).

[100] R. Sibille, N. Gauthier, H. Yan, M. Ciomaga Hatnean, J. Ollivier, B. Winn, U. Filges, G. Balakrishnan, M. Kenzelmann, N. Shannon, and T. Fennell, Nat. Phys. **14**, 711 (2018).

[101] T. Fennell, M. Kenzelmann, B. Roessli, M. K. Haas, and R. J. Cava, Phys. Rev. Lett. **109**, 017201 (2012).

[102] O. Benton, L. D. C. Jaubert, H. Yan, and N. Shannon, Nat. Comm. **7**, 11572 (2016).

[103] J. D. Thompson, P. A. McClarty, H. M. Rønnow, L. P. Regnault, A. Sorge, and M. J. Gingras, Phys. Rev. Lett. **106**, 187202 (2011).

[104] J. Robert, E. Lhotel, G. Remenyi, S. Sahling, I. Mirebeau, C. Decorse, B. Canals, and S. Petit, Phys. Rev. B **92**, 064425 (2015).

[105] J.-J. Wen, S. M. Koohpayeh, K. A. Ross, B. A. Trump, T. M. McQueen, K. Kimura, S. Nakatsuji, Y. Qiu, D. M. Pajerowski, J. R. D. Copley, and C. L. Broholm, Phys. Rev. Lett. **118**, 107206 (2017).

[106] J. G. Rau, R. Moessner, and P. A. McClarty, Phys. Rev. B **100**, 104423 (2019).

[107] L. D. Landau and E. M. Lifshitz, Phys. Z. Sowjetunion **8**, 153 (1935).

[108] I. I. Rabi, J. R. Zacharias, S. Millman, and P. Kusch, Phys. Rev. **53**, 318 (1938).

[109] P. H. Conlon and J. T. Chalker, Phys. Rev. B **81**, 224413 (2010).

[110] M. Taillefumier, J. Robert, C. L. Henley, R. Moessner, and B. Canals, Phys. Rev. B **90**, 064419 (2014).

[111] J. Robert, B. Canals, V. Simonet, and R. Ballou, Phys. Rev. Lett. **101**, 117207 (2008).

[112] S. Schnabel and D. P. Landau, Phys. Rev. B **86**, 014413 (2012).

[113] G. Baskaran, S. Mandal, and R. Shankar, Phys. Rev. Lett. **98**, 247201 (2007).

[114] G. Baskaran, D. Sen, and R. Shankar, Phys. Rev. B **78**, 115116 (2008).

[115] A. M. Samarakoon, A. Banerjee, S.-S. Zhang, Y. Kamiya, S. E. Nagler, D. A. Tennant, S.-H. Lee, and C. D. Batista, Phys. Rev. B **96**, 134408 (2017).

[116] A. M. Samarakoon, G. Wachtel, Y. Yamaji, D. A. Tennant, C. D. Batista, and Y. B. Kim, Phys. Rev. B **98**, 045121 (2018).

[117] C. Kittel, *Introduction to Solid State Physics*, 8th ed. (Wiley, New York, 2004).

[118] S.-S. Zhang, E. A. Ghioldi, Y. Kamiya, L. O. Manuel, A. E. Trumper, and C. D. Batista, Phys. Rev. B **100**, 104431 (2019).

[119] S. Onoda and Y. Tanaka, Phys. Rev. B **83**, 094411 (2011).

[120] J. Villain, J. Phys. France **35**, 27 (1974).

[121] G. T. Barkema and M. E. J. Newman, Phys. Rev. E **57**, 1155 (1998).

[122] D. J. Best and N. I. Fisher, Appl. Stat. **28**, 152 (1979).

[123] L. Devroye, *Non-Uniform Random Variate Generation* (Springer-Verlag, New York, 1986).

[124] M. Galassi, J. Davies, J. Theiler, B. Gough, G. Jungman, P. Alken, M. Booth, and F. Rossi, *GNU Scientific Library Reference Manual*, 3rd ed. (Network Theory, Bristol, 2009).

[125] P. Kovesi, arXiv:1509.03700 (2015).

[126] M. Frigo and S. Johnson, Proc. IEEE **93**, 216 (2005).

[127] S. W. de Leeuw, J. W. Perram, and E. R. Smith, Proc. R. Soc. London A **373**, 27 (1980).

[128] J. G. Rau and M. J. Gingras, Annu. Rev. Condens. Matter Phys. **10**, 357 (2019).

[129] G. Chen, Phys. Rev. B **94**, 205107 (2016).

[130] G. Chen, Phys. Rev. B **96**, 195127 (2017).

[131] R. Moessner and K. S. Raman, in *Introduction to Frustrated Magnetism*, Springer Series in Solid-State Sciences, Vol. 164, edited by C. Lacroix, P. Mendels, and F. Mila (Springer-Verlag, Berlin, 2011) Chap. 17.

[132] C. L. Henley, J. Stat. Phys. **89**, 483 (1997).

[133] P. Fendley, R. Moessner, and S. L. Sondhi, Phys. Rev. B **66**, 214513 (2002).

[134] R. Moessner and S. L. Sondhi, Phys. Rev. Lett. **86**, 1881 (2001).

[135] C. L. Henley, J. Phys. Condens. Matter **16**, S891 (2004).

[136] A. M. Läuchli, S. Capponi, and F. F. Assaad, J. Stat. Mech. **2008**, P01010 (2008).

[137] D. A. Ivanov, Phys. Rev. B **70**, 094430 (2004).

[138] C. Zeng and C. L. Henley, Phys. Rev. B **55**, 14935 (1997).

[139] P. M. Chaikin and T. C. Lubensky, *Principles of Condensed Matter Physics* (Cambridge University Press, 1995).

[140] R. Moessner and L. Sondhi, Phys. Rev. B **68**, 184512 (2003).

[141] A. Ioselevich, D. A. Ivanov, and M. V. Feigelman, Phys. Rev. B **66**, 174405 (2002).

[142] G. Misguich and F. Mila, Phys. Rev. B **77**, 134421 (2008).

[143] A. Ralko, M. Ferrero, F. Becca, D. Ivanov, and F. Mila, Phys. Rev. B **74**, 134301 (2006).

[144] G. Goldstein, C. Chamon, and C. Castelnovo, Phys. Rev. B **95**, 174511 (2017).

[145] A. Altland and B. D. Simons, *Condensed Matter Field Theory*, 2nd ed. (Cambridge University Press, Cambridge, 2010).

[146] Y. Tang, A. W. Sandvik, and C. L. Henley, Phys. Rev. B **84**, 174427 (2011).

[147] J. Villain, J. Phys. France **36**, 581 (1975).

[148] A. M. Polyakov, Phys. Lett. B **72**, 477 (1978).

[149] S. Sachdev, Phys. Rev. B **40**, 5204 (1989).

[150] P. Leung, K. Chiu, and K. J. Runge, Phys. Rev. B **54**, 12938 (1996).

[151] O. F. Syljuåsen, Phys. Rev. B **73**, 245105 (2006).

[152] A. Ralko, D. Poilblanc, and R. Moessner, Phys. Rev. Lett. **100**, 037201 (2008).

[153] D. Banerjee, M. Bögli, C. P. Hofmann, F. J. Jiang, P. Widmer, and U. J. Wiese, Phys. Rev. B **90**, 245143 (2014).

[154] D. Banerjee, M. Bögli, C. P. Hofmann, F. J. Jiang, P. Widmer, and U. J. Wiese, Phys. Rev. B **94**, 115120 (2016).

[155] T. Oakes, S. Powell, C. Castelnovo, A. Lamacraft, and J. P. Garrahan, Phys. Rev. B **98**, 064302 (2018).

[156] J. Herzog-Arbeitman, S. Mantilla, and I. Sodemann, Phys. Rev. B **99**, 245108 (2019).

[157] T. M. Schlittler, R. Mosseri, and T. Barthel, Phys. Rev. B **96**, 195142 (2017).

[158] O. Sikora, N. Shannon, F. Pollmann, K. Penc, and P. Fulde, Phys. Rev. B **84**, 115129 (2011).

[159] O. Sikora, F. Pollmann, N. Shannon, K. Penc, and P. Fulde, Phys. Rev. Lett. **103**, 247001 (2009).

[160] R. Moessner, S. L. Sondhi, and P. Chandra, Phys. Rev. B **64**, 144416 (2001).

[161] A. Ralko, M. Ferrero, F. Becca, D. Ivanov, and F. Mila, Phys. Rev. B **71**, 224109 (2005).

[162] E. Fradkin and S. H. Shenker, Phys. Rev. D **19**, 3682 (1979).

[163] G. Carleo, J. I. Cirac, K. Cranmer, L. Daudet, M. Schuld, N. Tishby, L. Vogt-Maranto, and L. Zdeborová, Rev. Mod. Phys. **91**, 045002 (2019).

[164] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning* (MIT Press, Cambridge, 2016).

[165] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. Brandao, D. A. Buell, B. Burkett, Y. Chen, Z. Chen, B. Chiaro, R. Collins, W. Courtney, A. Dunsworth, E. Farhi, B. Foxen, A. Fowler, C. Gidney, M. Giustina, R. Graff, K. Guerin, S. Habegger, M. P. Harrigan, M. J. Hartmann, A. Ho, M. Hoffmann, T. Huang, T. S. Humble, S. V. Isakov, E. Jeffrey, Z. Jiang, D. Kafri, K. Kechedzhi, J. Kelly, P. V. Klimov, S. Knysh, A. Korotkov, F. Kostritsa, D. Landhuis, M. Lindmark, E. Lucero, D. Lyakh, S. Mandrà, J. R. McClean, M. McEwen, A. Megrant, X. Mi, K. Michielsen, M. Mohseni, J. Mutus, O. Naaman, M. Neeley, C. Neill, M. Y. Niu, E. Ostby, A. Petukhov, J. C. Platt, C. Quintana, E. G. Rieffel, P. Roushan, N. C. Rubin, D. Sank, K. J. Satzinger, V. Smelyanskiy, K. J. Sung, M. D. Trevithick, A. Vainsencher, B. Villalonga, T. White, Z. J. Yao, P. Yeh, A. Zalcman, H. Neven, and J. M. Martinis, Nature **574**, 505 (2019).

[166] R. Orús, Nat. Rev. Phys. **1**, 538 (2019).

[167] C.-Y. Park and M. J. Kastoryano, Phys. Rev. Research **2**, 023232 (2020).

[168] S. Sorella, Phys. Rev. B **64**, 024512 (2001).

[169] Y. Iqbal, F. Becca, and D. Poilblanc, Phys. Rev. B **84**, 020407 (2011).

[170] Y. Iqbal, F. Becca, S. Sorella, and D. Poilblanc, Phys. Rev. B **87**, 060405 (2013).

[171] R. Schaffer, S. Bhattacharjee, and Y. B. Kim, Phys. Rev. B **86**, 224417 (2012).

[172] T. Li and F. Yang, Phys. Rev. B **81**, 214509 (2010).

[173] Y. Iqbal, D. Poilblanc, and F. Becca, Phys. Rev. B **89**, 020407 (2014).

[174] F. Ferrari and F. Becca, Phys. Rev. B **102**, 014417 (2020).

[175] Y. Nomura and M. Imada, arXiv:2005.14142 (2020).

[176] D. A. Huse and V. Elser, Phys. Rev. Lett. **60**, 2531 (1988).

[177] W. Marshall and R. E. Peierls, Proc. R. Soc. London A **232**, 48 (1955).

[178] F. Mezzacapo, N. Schuch, M. Boninsegni, and J. I. Cirac, New J. Phys. **11**, 083026 (2009).

[179] J. Thibaut, T. Roscilde, and F. Mezzacapo, Phys. Rev. B **100**, 155148 (2019).

[180] A. N. Rubtsov, V. V. Savkin, and A. I. Lichtenstein, Phys. Rev. B **72**, 035122 (2005).

[181] E. Gull, A. J. Millis, A. I. Lichtenstein, A. N. Rubtsov, M. Troyer, and P. Werner, Rev. Mod. Phys. **83**, 349 (2011).

[182] A. W. Sandvik, Phys. Rev. B **57**, 10287 (1998).

[183] D. M. Ceperley, in *Monte Carlo and Molecular Dynamics of Condensed Matter Systems*, edited by K. Binder and G. Ciccotti (Società Italiana di Fisica, Bologna, 1996).

[184] P. Henelius and A. W. Sandvik, Phys. Rev. B **62**, 1102 (2000).

[185] M. Troyer and U. J. Wiese, Phys. Rev. Lett. **94**, 170201 (2005).

[186] Z.-X. Li and H. Yao, Annu. Rev. Condens. Matter Phys. **10**, 337 (2019).

[187] R. R. Dos Santos, Braz. J. Phys. **33**, 36 (2003).

[188] S. Sorella, S. Baroni, R. Car, and M. Parrinello, Europhys. Lett. **8**, 663 (1989).

[189] W. M. Foulkes, L. Mitas, R. J. Needs, and G. Rajagopal, Rev. Mod. Phys. **73**, 33 (2001).

[190] F. G. Frobenius, Sitzungsber. K. Preuss. Akad. Wiss. **26**, 456 (1912).

[191] C. D. Meyer, *Matrix Analysis and Applied Linear Algebra* (SIAM, 2000).

[192] B. Vanhecke, J. Colbois, L. Vanderstraeten, F. Mila, and F. Verstraete, Phys. Rev. Research **3**, 013041 (2021).

[193] J. B. Anderson, J. Chem. Phys. **63**, 1499 (1975).

[194] J. B. Anderson, J. Chem. Phys. **65**, 4121 (1976).

[195] P. J. Reynolds, D. M. Ceperley, B. J. Alder, and W. A. Lester, J. Chem. Phys. **77**, 5593 (1982).

[196] J. Richter, N. B. Ivanov, and K. Retzlaff, Europhys. Lett. **25**, 545 (1994).

[197] A. Voigt, J. Richter, and N. B. Ivanov, Physica A **245**, 269 (1997).

[198] T. Westerhout, N. Astrakhantsev, K. S. Tikhonov, M. I. Katsnelson, and A. A. Bagrov, Nat. Comm. **11**, 1593 (2020).

[199] J. Schmidhuber, Neural Netw. **61**, 85 (2015).

[200] A. G. Ivakhnenko, IEEE Trans. Syst. Man Cybern. **1**, 364 (1971).

[201] A. Krizhevsky, I. Sutskever, and G. E. Hinton, Commun. ACM **60**, 84 (2017).

[202] G. Cybenko, Math. Control Signals Syst. **2**, 303 (1989).

[203] G. Carleo and M. Troyer, Science **355**, 602 (2017).

[204] Z. Cai and J. Liu, Phys. Rev. B **97**, 035116 (2018).

[205] K. Choo, T. Neupert, and G. Carleo, Phys. Rev. B **100**, 125124 (2019).

[206] J. Han, L. Zhang, and W. E, J. Comput. Phys. **399**, 108929 (2019).

[207] D. Pfau, J. S. Spencer, A. G. D. G. Matthews, and W. M. C. Foulkes, Phys. Rev. Research **2**, 033429 (2020).

[208] J. Hermann, Z. Schätzle, and F. Noé, Nat. Chem. **12**, 891 (2020).

[209] K. Choo, G. Carleo, N. Regnault, and T. Neupert, Phys. Rev. Lett. **121**, 167204 (2018).

[210] O. Sharir, Y. Levine, N. Wies, G. Carleo, and A. Shashua, Phys. Rev. Lett. **124**, 020503 (2020).

[211] M. Hibat-Allah, M. Ganahl, L. E. Hayward, R. G. Melko, and J. Carrasquilla, Phys. Rev. Research **2**, 023358 (2020).

[212] C. Roth, arXiv:2003.06228 (2020).

[213] A. Glielmo, Y. Rath, G. Csányi, A. De Vita, and G. H. Booth, Phys. Rev. X **10**, 041026 (2020).

[214] K. Choo, A. Mezzacapo, and G. Carleo, Nat. Comm. **11**, 2368 (2020).

[215] X. Liang, W.-Y. Liu, P.-Z. Lin, G.-C. Guo, Y.-S. Zhang, and L. He, Phys. Rev. B **98**, 104426 (2018).

[216] F. Ferrari, F. Becca, and J. Carrasquilla, Phys. Rev. B **100**, 125131 (2019).

[217] Y. Nomura, J. Phys. Soc. Jpn. **89**, 054706 (2020).

[218] Y. Levine, O. Sharir, N. Cohen, and A. Shashua, Phys. Rev. Lett. **122**, 065301 (2019).

[219] R. Kaubruegger, L. Pastori, and J. C. Budich, Phys. Rev. B **97**, 195136 (2018).

[220] R. G. Melko, G. Carleo, J. Carrasquilla, and J. I. Cirac, Nat. Phys. **15**, 887 (2019).

[221] P. Smolensky, in *Parallel Distributed Processing: Explorations in the Microstructure of Cognition, Volume 1: Foundations.*, edited by D. E. Rumelhart and J. L. McLelland (MIT Press, Cambridge, 1986) Chap. 6, pp. 194–281.

[222] N. Le Roux and Y. Bengio, Neural Comput. **20**, 1631 (2008).

[223] S. Geman and D. Geman, IEEE Trans. Pattern Anal. Mach. Intell. **PAMI-6**, 721 (1984).

[224] A. Graves, *Supervised Sequence Labelling with Recurrent Neural Networks*, 1st ed., Studies in Computational Intelligence, Vol. 385 (Springer-Verlag, Berlin, 2012).

[225] Y. Bengio, P. Lamblin, D. Popovici, and H. Larochelle, Adv. Neural Inf. Process. Syst. **19**, 153 (2007).

[226] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, Nature **323**, 533 (1986).

[227] S. d'Ascoli, L. Sagun, J. Bruna, and G. Biroli, arXiv:1906.06766 (2019).

[228] G. Carleo, K. Choo, D. Hofmann, J. E. Smith, T. Westerhout, F. Alet, E. J. Davis, S. Efthymiou, I. Glasser, S.-H. Lin, M. Mauri, G. Mazzola, C. B. Mendl, E. P. L. van Nieuwenburg, O. O'Reilly, H. Théveniaut, G. Torlai, F. Vicentini, and A. Wietek, SoftwareX **10**, 100311 (2019).

[229] D. Hendry and A. E. Feiguin, Phys. Rev. B **100**, 245123 (2019).

[230] G. Torlai, J. Carrasquilla, M. T. Fishman, R. G. Melko, and M. P. A. Fisher, Phys. Rev. Research **2**, 032060(R) (2020).

[231] X. Glorot and Y. Bengio, Proc. Mach. Learn. Res. **9**, 249 (2010).

[232] K. He, X. Zhang, S. Ren, and J. Sun, arXiv:1502.01852 (2015).

[233] L. Capriotti, F. Becca, A. Parola, and S. Sorella, Phys. Rev. Lett. **87**, 972011 (2001).

[234] H. C. Jiang, H. Yao, and L. Balents, Phys. Rev. B **86**, 024424 (2012).

[235] T. Li, F. Becca, W. Hu, and S. Sorella, Phys. Rev. B **86**, 075111 (2012).

[236] W.-J. Hu, F. Becca, A. Parola, and S. Sorella, Phys. Rev. B **88**, 060402(R) (2013).

[237] S.-S. Gong, W. Zhu, D. N. Sheng, O. I. Motrunich, and M. P. A. Fisher, Phys. Rev. Lett. **113**, 027201 (2014).

[238] S.-S. Gong, W. Zhu, and D. N. Sheng, Scientific Reports **4**, 1 (2014).

[239] T. Vieijra, C. Casert, J. Nys, W. De Neve, J. Haegeman, J. Ryckebusch, and F. Verstraete, Phys. Rev. Lett. **124**, 097201 (2020).

[240] C. Baldassi, F. Pittorino, and R. Zecchina, Proc. Natl. Acad. Sci. **117**, 161 (2020).

[241] D. Sehayek, A. Golubeva, M. S. Albergo, B. Kulchytskyy, G. Torlai, and R. G. Melko, Phys. Rev. B **100**, 195125 (2019).

[242] H. Théveniaut and F. Alet, "Combining RBM and reptation QMC: application to a 2d bosonic model," (2020), paper presented at the Machine Learning for Quantum Simulation conference.

[243] C. Roth and A. H. MacDonald, arXiv:2104.05085 (2021).

[244] T. S. Cohen and M. Welling, arXiv:1602.07576 (2016).

[245] H. W. Lin, M. Tegmark, and D. Rolnick, J. Stat. Phys. **168**, 1223 (2017).

[246] G. Valle-Pérez, C. Q. Camargo, and A. A. Louis, arXiv:1805.08522 (2018).

[247] G. De Palma, B. T. Kiani, and S. Lloyd, arXiv:1812.10156 (2018).

[248] A. Trask, F. Hill, S. Reed, J. Rae, C. Dyer, and P. Blunsom, arXiv:1808.00508 (2018).

[249] A. van den Oord, N. Kalchbrenner, O. Vinyals, L. Espeholt, A. Graves, and K. Kavukcuoglu, Adv. Neural Inf. Process. Syst. **29**, 4797 (2016).

[250] R. P. Feynman and M. Cohen, Phys. Rev. **102**, 1189 (1956).

[251] D. Luo and B. K. Clark, Phys. Rev. Lett. **122**, 226401 (2019).

[252] V. Baltz, A. Manchon, M. Tsoi, T. Moriyama, T. Ono, and Y. Tserkovnyak, Rev. Mod. Phys. **90**, 015005 (2018).

[253] W. Witczak-Krempa, G. Chen, Y. B. Kim, and L. Balents, Annu. Rev. Condens. Matter Phys. **5**, 57 (2014).

[254] J. G. Rau, E. K.-H. Lee, and H.-Y. Kee, Annu. Rev. Condens. Matter Phys. **7**, 195 (2016).

[255] B. Yan and S. C. Zhang, Rep. Prog. Phys. **75**, 096501 (2012).

[256] G. Jackeli and G. Khaliullin, Phys. Rev. Lett. **102**, 017205 (2009).

[257] Y. Kasahara, T. Ohnishi, Y. Mizukami, O. Tanaka, S. Ma, K. Sugii, N. Kurita, H. Tanaka, J. Nasu, Y. Motome, T. Shibauchi, and Y. Matsuda, Nature **559**, 227 (2018).

[258] K. Matsuhira, M. Wakeshima, Y. Hinatsu, and S. Takagi, J. Phys. Soc. Jpn. **80**, 094701 (2011).

[259] B. Canals and C. Lacroix, Phys. Rev. Lett. **80**, 2933 (1998).

[260] B. Canals and C. Lacroix, Phys. Rev. B **61**, 1149 (2000).

[261] M. Elhajal, B. Canals, R. Sunyer, and C. Lacroix, Phys. Rev. B **71**, 094420 (2005).

[262] J. J. Ishikawa, E. C. T. O'Farrell, and S. Nakatsuji, Phys. Rev. B **85**, 245109 (2012).

[263] K. Ueda, J. Fujioka, and Y. Tokura, Phys. Rev. B **93**, 245120 (2016).

[264] X. Wan, A. M. Turner, A. Vishwanath, and S. Y. Savrasov, Phys. Rev. B **83**, 205101 (2011).

[265] Y. Yamaji and M. Imada, Phys. Rev. X **4**, 021035 (2014).

[266] K. Wang, B. Xu, C. W. Rischau, N. Bachar, B. Michon, J. Teyssier, Y. Qiu, T. Ohtsuki, B. Cheng, N. P. Armitage, S. Nakatsuji, and D. van der Marel, Nat. Phys. **16**, 1194 (2020).

[267] H. Jacobsen, C. D. Dashwood, E. Lhotel, D. Khalyavin, P. Manuel, R. Stewart, D. Prabhakaran, D. F. McMorrow, and A. T. Boothroyd, Phys. Rev. B **101**, 104404 (2020).

[268] V. Cathelin, E. Lefrançois, J. Robert, P. C. Guruciaga, C. Paulsen, D. Prabhakaran, P. Lejay, F. Damay, J. Ollivier, B. Fåk, L. C. Chapon, R. Ballou, V. Simonet, P. C. W. Holdsworth, and E. Lhotel, Phys. Rev. Research **2**, 032073(R) (2020).

[269] J. N. Millican, R. T. Macaluso, S. Nakatsuji, Y. Machida, Y. Maeno, and J. Y. Chan, Mater. Res. Bull. **42**, 928 (2007).

[270] M. J. Harris, S. T. Bramwell, P. C. W. Holdsworth, and J. D. Champion, Phys. Rev. Lett. **81**, 4496 (1998).

[271] A. Petrenko, R. Lees, and G. Balakrishnan, Phys. Rev. B **68**, 012406 (2003).

[272] L. D. C. Jaubert and P. C. W. Holdsworth, Nat. Phys. **5**, 258 (2009).

[273] D. C. Jiles, *Introduction to Magnetism and Magnetic Materials*, 2nd ed. (Chapman & Hall, London, 1998).

[274] M. Saito, R. Higashinaka, and Y. Maeno, Phys. Rev. B **72**, 144422 (2005).

[275] A. Sarkar and S. Mukhopadhyay, Phys. Rev. B **90**, 165129 (2014).

[276] L. Lin, Y. L. Xie, J.-J. Wen, S. Dong, Z. B. Yan, and J.-M. Liu, New J. Phys. **17**, 123018 (2015).

[277] D. Khomskii, Nat. Comm. **3**, 904 (2012).

[278] T. H. Arima, J. Phys. Soc. Jpn. **82**, 013705 (2013).

[279] Z. Tian, Y. Kohama, T. Tomita, H. Ishizuka, T. H. Hsieh, J. J. Ishikawa, K. Kindo, L. Balents, and S. Nakatsuji, Nat. Phys. **12**, 134 (2016).

[280] N. W. Ashcroft and N. D. Mermin, *Solid state physics* (Holt, Rinehart & Winston, New York, 1976).

[281] I. A. Ryzhkin, Sov. Phys. JETP **101**, 481 (2005).

[282] V. Kaiser, S. T. Bramwell, P. C. W. Holdsworth, and R. Moessner, Nat. Mater. **12**, 1033 (2013).

[283] V. Kaiser, S. T. Bramwell, P. C. W. Holdsworth, and R. Moessner, Phys. Rev. Lett. **115**, 037201 (2015).

[284] R. Berger, Mem. Am. Math. Soc. **66**, 1 (1966).

[285] R. Penrose, Bull. Inst. Math. Applic. **10**, 266 (1974).

[286] F. P. M. Beenker, *Algebraic theory of non-periodic tilings of the plane by two simple building blocks: a square and a rhombus* (Eindhoven University of Technology, 1982).

[287] D. Shechtman and I. A. Blech, Metall. Trans. A. **16A**, 1005 (1985).

[288] P. G. Harper, Proc. Phys. Soc. A **68**, 874 (1955).

[289] M. Ya. Azbel', Sov. Phys. JETP **19**, 634 (1964).

[290] D. R. Hofstadter, Phys. Rev. B **14**, 2239 (1976).

[291] G. Roati, C. D'Errico, L. Fallani, M. Fattori, C. Fort, M. Zaccanti, G. Modugno, M. Modugno, and M. Inguscio, Nature **453**, 895 (2008).

[292] B. Deissler, M. Zaccanti, G. Roati, C. D'Errico, M. Fattori, M. Modugno, G. Modugno, and M. Inguscio, Nat. Phys. **6**, 354 (2010).

[293] L. Sanchez-Palencia and M. Lewenstein, Nat. Phys. **6**, 87 (2010).

[294] M. Schreiber, S. S. Hodgman, P. Bordia, H. P. Lüschen, M. H. Fischer, R. Vosk, E. Altman, U. Schneider, and I. Bloch, Science **349**, 842 (2015).

[295] P. Bordia, H. P. Lüschen, S. S. Hodgman, M. Schreiber, I. Bloch, and U. Schneider, Phys. Rev. Lett. **116**, 140401 (2016).

[296] P. Bordia, H. P. Lüschen, S. Scherg, S. Gopalakrishnan, M. Knap, U. Schneider, and I. Bloch, Phys. Rev. X **7**, 41047 (2017).

[297] H. P. Lüschen, P. Bordia, S. Scherg, F. Alet, E. Altman, U. Schneider, and I. Bloch, Phys. Rev. Lett. **119**, 260401 (2017).

[298] K. Viebahn, M. Sbroscia, E. Carter, J.-C. Yu, and U. Schneider, Phys. Rev. Lett. **122**, 110404 (2019).

[299] T. A. Corcovilos and J. Mittal, Appl. Opt. **58**, 2256 (2019).

[300] T. Kohlert, S. Scherg, X. Li, H. P. Lüschen, S. Das Sarma, I. Bloch, and M. Aidelsburger, Phys. Rev. Lett. **122**, 170403 (2019).

[301] T. S. Cubitt, D. Perez-Garcia, and M. M. Wolf, Nature **528**, 207 (2015).

[302] P. W. Anderson, Phys. Rev. **109**, 1492 (1958).

[303] M. Kohmoto, L. P. Kadanoff, and C. Tang, Phys. Rev. Lett. **50**, 1870 (1983).

[304] S. Ostlund, R. Pandit, D. Rand, H. J. Schellnhuber, and E. D. Siggia, Phys. Rev. Lett. **50**, 1873 (1983).

[305] J. Q. You, J. R. Yan, T. Xie, X. Zeng, and J. X. Zhong, J. Phys. Condens. Matter **3**, 7255 (1991).

[306] Q. Niu and F. Nori, Phys. Rev. Lett. **57**, 2057 (1986).

[307] Q. Niu and F. Nori, Phys. Rev. B **42**, 10329 (1990).

[308] S. Aubry and G. André, Ann. Israel Phys. Soc. **3**, 133 (1980).

[309] J. H. Han, D. J. Thouless, H. Hiramoto, and M. Kohmoto, Phys. Rev. B **50**, 11365 (1994).

[310] F. Liu, S. Ghosh, and Y. D. Chong, Phys. Rev. B **91**, 14108 (2015).

[311] T. Devakul and D. A. Huse, Phys. Rev. B **96**, 214201 (2017).

[312] J. Sutradhar, S. Mukerjee, R. Pandit, and S. Banerjee, Phys. Rev. B **99**, 224204 (2019).

[313] T. Vojta, J. Low Temp. Phys. **161**, 299 (2010).

[314] S. Gopalakrishnan, K. Agarwal, E. A. Demler, D. A. Huse, and M. Knap, Phys. Rev. B **93**, 134206 (2016).

[315] V. Khemani, D. N. Sheng, and D. A. Huse, Phys. Rev. Lett. **119**, 075702 (2017).

[316] W. De Roeck and F. Huveneers, Phys. Rev. B **95**, 155129 (2017).

[317] Y. E. Kraus, Y. Lahini, Z. Ringel, M. Verbin, and O. Zilberberg, Phys. Rev. Lett. **109**, 106402 (2012).

[318] M. Senechal, *Quasicrystals and geometry* (Cambridge University Press, Cambridge, 1995).

[319] Y. E. Kraus, Z. Ringel, and O. Zilberberg, Phys. Rev. Lett. **111**, 226401 (2013).

[320] O. Zilberberg, S. Huang, J. Guglielmon, M. Wang, K. P. Chen, Y. E. Kraus, and M. C. Rechtsman, Nature **553**, 59 (2018).

[321] I. M. Suslov, Sov. Phys. JETP **56**, 612 (1982).

[322] M. Wilkinson, Proc. R. Soc. London A **391**, 305 (1984).

[323] C. M. Soukoulis and E. N. Economou, Phys. Rev. Lett. **48**, 1043 (1982).

[324] M. Ya. Azbel', Phys. Rev. Lett. **43**, 1954 (1979).

[325] J. A. Ketoja and I. I. Satija, Phys. Lett. A **194**, 64 (1994).

[326] J. A. Ketoja and I. I. Satija, Physica A **219**, 212 (1995).

[327] J. Hermisson, U. Grimm, and M. Baake, J. Phys. A **30**, 7315 (1997).

[328] B. D. Mestel and A. H. Osbaldestin, J. Phys. A **37**, 9071 (2004).

[329] D. Karevski and L. Turban, J. Phys. A **29**, 3461 (1996).

[330] M. Nauenberg, J. Phys. A **8**, 925 (1975).

[331] D. Bessis, J. S. Geronimo, and P. Moussa, J. Physique Lett. **44**, 977 (1983).

[332] B. Doucot, W. Wang, J. Chaussy, B. Pannetier, R. Rammal, A. Vareille, and D. Henry, Phys. Rev. Lett. **57**, 1235 (1986).

[333] B. Derrida, C. Itzykson, and J. M. Luck, Comm. Math. Phys. **94**, 115 (1984).

[334] N. Mott, J. Phys. C **20**, 3075 (1987).

[335] S. Ya. Jitomirskaya, Ann. Math. **150**, 1159 (1999).

[336] D. J. Thouless, Phys. Rev. B **28**, 4272 (1983).

[337] C. Tang and M. Kohmoto, Phys. Rev. B **34**, 2041 (1986).

[338] R. B. Stinchcombe and S. C. Bell, J. Phys. A **20**, L739 (1987).

[339] S. C. Bell and R. B. Stinchcombe, J. Phys. A **22**, 717 (1989).

[340] M. Kohmoto and J. R. Banavar, Phys. Rev. B **34**, 563 (1986).

[341] Y. Hatsugai and M. Kohmoto, Phys. Rev. B **42**, 8282 (1990).

[342] A. Avila and S. Jitomirskaya, Ann. Math. **170**, 303 (2009).

[343] J. Bellissard, A. Bovier, and J.-M. Ghez, Rev. Math. Phys. **04**, 1 (1992).

[344] D. J. Thouless, M. Kohmoto, M. P. Nightingale, and M. den Nijs, Phys. Rev. Lett. **49**, 405 (1982).

[345] I. Dana, Phys. Rev. B **89**, 205111 (2014).

[346] S. Lang, *Introduction to Diophantine Approximations* (Springer-Verlag, 1995).

[347] L. D. Landau and E. M. Lifshitz, *Quantum Mechanics: Non-Relativistic Theory*, 3rd ed., Course of Theoretical Physics, Vol. 3 (Pergamon, Oxford, 1977).

[348] E. M. Lifshitz and L. P. Pitaevskii, *Statistical Physics, Part 2*, Course of Theoretical Physics, Vol. 9 (Pergamon, Oxford, 1980).

[349] E. H. Lieb, R. Seiringer, and J. Yngvason, Phys. Rev. B **66**, 134529 (2002).

[350] R. Roth and K. Burnett, Phys. Rev. A **68**, 23604 (2003).

[351] J. C. C. Cestari, A. Foerster, and M. A. Gusmão, Phys. Rev. A **82**, 63634 (2010).

[352] M. P. A. Fisher, P. B. Weichman, G. Grinstein, and D. S. Fisher, Phys. Rev. B **40**, 546 (1989).

[353] S. Sachdev, *Quantum Phase Transitions* (Cambridge University Press, 2000).

[354] V. K. Varma, C. de Mulatier, and M. Žnidarič, Phys. Rev. E **96**, 32130 (2017).

[355] H. Hiramoto and S. Abe, J. Phys. Soc. Jpn. **57**, 1365 (1988).

[356] H. Hiramoto and S. Abe, J. Phys. Soc. Jpn. **57**, 230 (1988).

[357] S. Abe and H. Hiramoto, Phys. Rev. A **36**, 5349 (1987).

[358] F. Piéchon, Phys. Rev. Lett. **76**, 4372 (1996).

[359] T. C. Halsey, M. H. Jensen, L. P. Kadanoff, I. Procaccia, and B. I. Shraiman, Phys. Rev. A **33**, 1141 (1986).

[360] M. Kohmoto, B. Sutherland, and C. Tang, Phys. Rev. B **35**, 1020 (1987).

[361] D. J. Thouless, Comm. Math. Phys. **127**, 187 (1990).

[362] Y. Last and M. Wilkinson, J. Phys. A **25**, 6123 (1992).

[363] Y. Tan, J. Phys. A **28**, 4163 (1995).

[364] M. Rossignolo and L. Dell'Anna, Phys. Rev. B **99**, 054211 (2019).

[365] A.-M. Guo, X. C. Xie, and Q.-f. Sun, Phys. Rev. B **89**, 075434 (2014).

[366] H.-j. Li, J.-p. Dou, and G. Huang, Phys. Rev. A **89**, 033843 (2014).

[367] A. Chandran and C. R. Laumann, Phys. Rev. X **7**, 031061 (2017).

[368] B. Huang and W. V. Liu, Phys. Rev. B **100**, 144202 (2019).

[369] T. Kariyado and A. Vishwanath, Phys. Rev. Research **1**, 033076 (2019).

[370] E. Abrahams, P. W. Anderson, D. C. Licciardello, and T. V. Ramakrishnan, Phys. Rev. Lett. **42**, 673 (1979).

[371] P. J. D. Crowley and A. Chandran, Phys. Rev. Research **2**, 033262 (2020).

[372] M. Sbroscia, K. Viebahn, E. Carter, J.-C. Yu, A. Gaunt, and U. Schneider, Phys. Rev. Lett. **125**, 200604 (2020).

[373] R.-J. Slager, personal communication.

[374] J. Schachenmayer, A. Pikovski, and A. M. Rey, Phys. Rev. X **5**, 011022 (2015).

[375] S. E. Begg, A. G. Green, and M. J. Bhaseen, J. Phys. A **53**, 50LT02 (2020).

[376] S. De Nicola, B. Doyon, and M. J. Bhaseen, J. Stat. Mech. **2020**, 13106 (2020).

[377] S. De Nicola, J. Stat. Mech. **2021**, 013101 (2021).

[378] V. Heine, *Group Theory in Quantum Mechanics*, International Series in Natural Philosophy, Vol. 91 (Pergamon, Oxford, 1960).

[379] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. J. Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey, I. Polat, Y. Feng, E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, and SciPy 1.0 Contributors, Nat. Methods **17**, 261 (2020).

[380] J. D. Jackson, *Classical electrodynamics*, 3rd ed. (Wiley, New York, 1998).